OFED+ Host Software Release 1.5.4

User Guide

IB0054606-02 A

Information furnished in this manual is believed to be accurate and reliable. However, QLogic Corporation assumes no responsibility for its use, nor for any infringements of patents or other rights of third parties which may result from its use. QLogic Corporation reserves the right to change product specifications at any time without notice. Applications described in this document for any of these products are for illustrative purposes only. QLogic Corporation makes no representation nor warranty that such applications are suitable for the specified use without further testing or modification. QLogic Corporation assumes no responsibility for any errors that may appear in this document.

Document Revision History	
Revision A, April 2012	
Changes	**Sections Affected**

Table of Contents

Preface

Intended Audience xv
Related Materials xv
Documentation Conventions xv
License Agreements xvi
Technical Support xvii
Training xvii
Contact Information xvii
Knowledge Database xviii

1 Introduction

How this Guide is Organized 1-1
Overview 1-2
Interoperability 1-3

2 Step-by-Step Cluster Setup and MPI Usage Checklists

Cluster Setup 2-1
Using MPI 2-2

3 InfiniBand® Cluster Setup and Administration

Introduction 3-1
Installed Layout 3-2
IB and OpenFabrics Driver Overview 3-3
IPoIB Network Interface Configuration 3-3
IPoIB Administration 3-5
Administering IPoIB 3-5
Stopping, Starting and Restarting the IPoIB Driver 3-5
Configuring IPoIB 3-5
Editing the IPoIB Configuration File 3-5
IB Bonding 3-6
Interface Configuration Scripts 3-6
Red Hat EL5 and EL6 3-7
SuSE Linux Enterprise Server (SLES) 10 and 11 3-8
Verify IB Bonding is Configured 3-9

Subnet Manager Configuration 3-10
QLogic Distributed Subnet Administration 3-12
Applications that use Distributed SA 3-12
Virtual Fabrics and the Distributed SA. 3-13
Configuring the Distributed SA 3-13
Default Configuration. 3-13
Multiple Virtual Fabrics Example 3-14
Virtual Fabrics with Overlapping Definitions 3-15
Distributed SA Configuration File 3-17
SID 3-18
ScanFrequency 3-18
LogFile 3-18
Dbg 3-19
Other Settings. 3-19
Changing the MTU Size 3-20
Managing the ib_qib Driver 3-21
Configure the ib_qib Driver State. 3-22
Start, Stop, or Restart ib_qib Driver 3-22
Unload the Driver/Modules Manually. 3-23
ib_qib Driver Filesystem 3-23
More Information on Configuring and Loading Drivers. 3-24
Performance Settings and Management Tips 3-24
Performance Tuning 3-25
Systems in General (With Either Intel or AMD CPUs) 3-25
AMD CPU Systems 3-28
AMD Interlagos CPU Systems 3-28
Intel CPU Systems 3-28
High Risk Tuning for Intel Harpertown CPUs 3-30
Additional Driver Module Parameter Tunings Available 3-31
Performance Tuning using `ipath_perf_tuning` Tool 3-34
OPTIONS 3-35
AUTOMATIC vs. INTERACTIVE MODE. 3-36
Affected Files 3-37
Homogeneous Nodes 3-37
Adapter and Other Settings. 3-38
Remove Unneeded Services. 3-39
Host Environment Setup for MPI 3-40

Configuring for ssh 3-40
Configuring ssh and sshd Using shosts.equiv 3-40
Configuring for ssh Using ssh-agent 3-43
Process Limitation with ssh 3-44
Checking Cluster and Software Status 3-44
ipath_control 3-44
iba_opp_query 3-45
ibstatus 3-46
ibv_devinfo 3-47
ipath_checkout 3-47

4 Running MPI on QLogic Adapters

Introduction 4-1
MPIs Packaged with QLogic OFED+ 4-1
Open MPI 4-1
Installation 4-2
Setup 4-2
Compiling Open MPI Applications 4-2
Create the mpihosts File 4-3
Running Open MPI Applications 4-3
Further Information on Open MPI 4-4
Configuring MPI Programs for Open MPI 4-5
To Use Another Compiler 4-5
Compiler and Linker Variables 4-7
Process Allocation 4-7
IB Hardware Contexts on the QDR IB Adapters 4-8
Enabling and Disabling Software Context Sharing 4-9
Restricting IB Hardware Contexts in a Batch Environment 4-10
Context Sharing Error Messages 4-11
Running in Shared Memory Mode 4-11
mpihosts File Details 4-12
Using Open MPI's mpirun 4-13
Console I/O in Open MPI Programs 4-14
Environment for Node Programs 4-15
Remote Execution 4-15
Exported Environment Variables 4-16
Setting MCA Parameters 4-17
Environment Variables 4-18
Job Blocking in Case of Temporary IB Link Failures 4-20
Open MPI and Hybrid MPI/OpenMP Applications 4-21

Debugging MPI Programs 4-22
MPI Errors 4-22
Using Debuggers 4-22

5 Using Other MPIs

Introduction 5-1
Installed Layout 5-2
Open MPI 5-3
MVAPICH 5-3
Compiling MVAPICH Applications 5-3
Running MVAPICH Applications 5-3
Further Information on MVAPICH 5-4
MVAPICH2 5-4
Compiling MVAPICH2 Applications 5-4
Running MVAPICH2 Applications 5-5
Further Information on MVAPICH2 5-5
Managing MVAPICH, and MVAPICH2 with the `mpi-selector` Utility 5-5
Platform MPI 8 5-6
Installation 5-6
Setup 5-6
Compiling Platform MPI 8 Applications 5-7
Running Platform MPI 8 Applications 5-7
More Information on Platform MPI 8 5-7
Intel MPI 5-7
Installation 5-8
Setup 5-8
Compiling Intel MPI Applications 5-10
Running Intel MPI Applications 5-10
Further Information on Intel MPI 5-11
Improving Performance of Other MPIs Over IB Verbs 5-12

6 SHMEM Description and Configuration

Overview 6-1
Interoperability 6-1
Installation 6-1
SHMEM Programs 6-3
Basic SHMEM Program 6-3
Compiling SHMEM Programs 6-4

Running SHMEM Programs 6-5
Using shmemrun 6-5
Running programs without using shmemrun 6-6
QLogic SHMEM Relationship with MPI 6-7
Slurm Integration 6-8
Full Integration 6-8
Two-step Integration 6-8
No Integration 6-9
Sizing Global Shared Memory 6-9
Progress Model 6-11
Active Progress 6-12
Passive Progress 6-12
Active versus Passive Progress 6-13
Environment Variables 6-13
Implementation Behavior 6-15
Application Programming Interface 6-17
SHMEM Benchmark Programs 6-27

7 Virtual Fabric support in PSM

Introduction 7-1
Virtual Fabric Support 7-2
Using SL and PKeys 7-2
Using Service ID 7-3
SL2VL mapping from the Fabric Manager 7-3
Verifying SL2VL tables on QLogic 7300 Series Adapters 7-4

8 Dispersive Routing

9 gPXE

gPXE Setup 9-1
Required Steps 9-2
Preparing the DHCP Server in Linux 9-2
Installing DHCP 9-3
Configuring DHCP 9-4
Netbooting Over IB 9-5
Prerequisites 9-5
Boot Server Setup 9-5
Steps on the gPXE Client 9-14
HTTP Boot Setup 9-14

A Benchmark Programs

Benchmark 1: Measuring MPI Latency Between Two Nodes A-1
Benchmark 2: Measuring MPI Bandwidth Between Two Nodes A-4
Benchmark 3: Messaging Rate Microbenchmarks A-6
OSU Multiple Bandwidth / Message Rate test (osu_mbw_mr) A-6
An Enhanced Multiple Bandwidth / Message Rate test (mpi_multibw) A-7

B SRP Configuration

SRP Configuration Overview B-1
Important Concepts B-1
QLogic SRP Configuration B-2
Stopping, Starting and Restarting the SRP Driver B-3
Specifying a Session B-3
Determining the values to use for the configuration B-6
Specifying an SRP Initiator Port of a Session by Card and Port Indexes B-8
Specifying an SRP Initiator Port of Session by Port GUID B-8
Specifying a SRP Target Port B-9
Specifying a SRP Target Port of a Session by IOCGUID B-10
Specifying a SRP Target Port of a Session by Profile String . . . B-10
Specifying an Adapter B-10
Restarting the SRP Module B-11
Configuring an Adapter with Multiple Sessions B-11
Configuring Fibre Channel Failover B-13
Failover Configuration File 1: Failing over from one SRP Initiator port to another B-14
Failover Configuration File 2: Failing over from a port on the VIO hardware card to another port on the VIO hardware card. B-15
Failover Configuration File 3: Failing over from a port on a VIO hardware card to a port on a different VIO hardware card within the same Virtual I/O chassis B-16
Failover Configuration File 4: Failing over from a port on a VIO hardware card to a port on a different VIO hardware card in a different Virtual I/O chassis B-17
Configuring Fibre Channel Load Balancing B-18
1 Adapter Port and 2 Ports on a Single VIO B-18
2 Adapter Ports and 2 Ports on a Single VIO Module B-19
Using the roundrobinmode Parameter B-20

Configuring SRP for Native IB Storage B-21
Notes B-23
Additional Details B-24
Troubleshooting B-24
OFED SRP Configuration. B-24

C Integration with a Batch Queuing System

Clean Termination of MPI Processes C-1
Clean-up PSM Shared Memory Files. C-2

D Troubleshooting

Using LEDs to Check the State of the Adapter D-1
BIOS Settings. D-2
Kernel and Initialization Issues. D-2
Driver Load Fails Due to Unsupported Kernel. D-3
Rebuild or Reinstall Drivers if Different Kernel Installed D-3
InfiniPath Interrupts Not Working. D-3
OpenFabrics Load Errors if ib_qib Driver Load Fails D-4
InfiniPath ib_qib Initialization Failure. D-5
MPI Job Failures Due to Initialization Problems D-6
OpenFabrics and InfiniPath Issues D-6
Stop Infinipath Services Before Stopping/Restarting InfiniPath D-6
Manual Shutdown or Restart May Hang if NFS in Use D-7
Load and Configure IPoIB Before Loading SDP D-7
Set $IBPATH for OpenFabrics Scripts D-7
SDP Module Not Loading D-7
ibsrpdm Command Hangs when Two Host Channel Adapters are Installed but Only Unit 1 is Connected to the Switch D-8
Outdated ipath_ether Configuration Setup Generates Error D-8
System Administration Troubleshooting. D-8
Broken Intermediate Link. D-9
Performance Issues D-9
Large Message Receive Side Bandwidth Varies with Socket Affinity on Opteron Systems D-9
Erratic Performance. D-10
Method 1. . . . . D-10
Method 2. . . . . D-10
Immediately change the processor affinity of an IRQ. D-11
Performance Warning if ib_qib Shares Interrupts with eth0 . . . . . D-12

Open MPI Troubleshooting D-12
Invalid Configuration Warning D-12

E ULP Troubleshooting

Troubleshooting VirtualNIC and VIO Hardware Issues E-1
Checking the logical connection between the IB Host and the VIO hardware E-1
Verify that the proper VirtualNIC driver is running E-2
Verifying that the qlgc_vnic.cfg file contains the correct information E-2
Verifying that the host can communicate with the I/O Controllers (IOCs) of the VIO hardware E-3
Checking the interface definitions on the host E-6
Interface does not show up in output of 'ifconfig' E-6
Verify the physical connection between the VIO hardware and the Ethernet network E-7
Troubleshooting SRP Issues E-9
ib_qlgc_srp_stats showing session in disconnected state E-9
Session in 'Connection Rejected' state E-11
Attempts to read or write to disk are unsuccessful E-14
Four sessions in a round-robin configuration are active E-15
Which port does a port GUID refer to? E-16
How does the user find a HCA port GUID? E-17
Need to determine the SRP driver version E-19

F Write Combining

Introduction F-1
PAT and Write Combining F-1
MTRR Mapping and Write Combining F-2
Edit BIOS Settings to Fix MTRR Issues F-2
Use the ipath_mtrr Script to Fix MTRR Issues F-2
Verify Write Combining is Working F-3

G Commands and Files

Check Cluster Homogeneity with ipath_checkout G-1
Restarting InfiniPath G-2
Summary and Descriptions of Commands G-2
dmesg G-4
iba_opp_query G-4
iba_hca_rev G-9
iba_manage_switch G-19

iba_packet_capture G-21
ibhosts G-22
ibstatus G-22
ibtracert G-23
ibv_devinfo G-24
ident G-24
ipath_checkout G-25
ipath_control G-27
ipath_mtrr G-28
ipath_pkt_test G-29
ipathstats G-30
lsmod G-30
modprobe G-30
mpirun G-31
mpi_stress G-31
rpm G-32
strings G-32
Common Tasks and Commands G-32
Summary and Descriptions of Useful Files G-34
boardversion G-34
status_str G-35
version G-36
Summary of Configuration Files G-36

H Recommended Reading

References for MPI H-1
Books for Learning MPI Programming H-1
Reference and Source for SLURM H-1
InfiniBand® H-1
OpenFabrics H-2
Clusters H-2
Networking H-2
Rocks H-2
Other Software Packages H-2

List of Figures

3-1 QLogic OFED+ Software Structure 3-1
3-2 Distributed SA Default Configuration 3-13
3-3 Distributed SA Multiple Virtual Fabrics Example 3-14
3-4 Distributed SA Multiple Virtual Fabrics Configured Example 3-15
3-5 Virtual Fabrics with Overlapping Definitions. 3-15
3-6 Virtual Fabrics with PSM_MPI Virtual Fabric Enabled 3-16
3-7 Virtual Fabrics with all SIDs assigned to PSM_MPI Virtual Fabric. 3-16
3-8 Virtual Fabrics with Unique Numeric Indexes 3-17

List of Tables

3-1 ibmtu Values . . . 3-20
3-2 **krcvqs** Parameter Settings . . . 3-27
3-3 Checks Preformed by ipath_perf_tuning Tool . . . 3-34
3-4 ipath_perf_tuning Tool Options . . . 3-35
3-5 Test Execution Modes . . . 3-36
4-1 Open MPI Wrapper Scripts . . . 4-2
4-2 Command Line Options for Scripts . . . 4-3
4-3 Intel . . . 4-6
4-4 Portland Group (PGI) . . . 4-6
4-5 Available Hardware and Software Contexts . . . 4-8
4-6 Environment Variables Relevant for any PSM . . . 4-18
4-7 Environment Variables Relevant for Open MPI . . . 4-20
5-1 Other Supported MPI Implementations . . . 5-1
5-2 MVAPICH Wrapper Scripts . . . 5-3
5-3 MVAPICH Wrapper Scripts . . . 5-4
5-4 Platform MPI 8 Wrapper Scripts . . . 5-7
5-5 Intel MPI Wrapper Scripts . . . 5-10
6-1 SHMEM Run Time Library Environment Variables . . . 6-13
6-2 **shmemrun** Environment Variables . . . 6-15
6-3 SHMEM Application Programming Interface Calls . . . 6-18
6-4 QLogic SHMEM micro-benchmarks options . . . 6-27
6-5 QLogic SHMEM random access benchmark options . . . 6-28
6-6 QLogic SHMEM all-to-all benchmark options . . . 6-29
6-7 QLogic SHMEM barrier benchmark options . . . 6-30
6-8 QLogic SHMEM reduce benchmark options . . . 6-31
D-1 LED Link and Data Indicators . . . D-1
G-1 Useful Programs . . . G-2
G-2 ipath_checkout Options . . . G-26
G-3 Common Tasks and Commands Summary . . . G-33
G-4 Useful Files . . . G-34
G-5 status_str File Contents . . . G-35
G-6 Status—Other Files . . . G-36
G-7 Configuration Files . . . G-37

Preface

The *QLogic OFED+ Host Software User Guide* shows end users how to use the installed software to setup the fabric. End users include both the cluster administrator and the Message-Passing Interface (MPI) application programmers, who have different but overlapping interests in the details of the technology.

For specific instructions about installing the QLogic QLE7340, QLE7342, QMH7342, and QME7342 PCI Express® (PCIe®) adapters see the *QLogic InfiniBand®* Adapter Hardware Installation Guide, and the initial installation of the Fabric Software, see the *QLogic InfiniBand®* Fabric Software Installation Guide.

Intended Audience

This guide is intended for end users responsible for administration of a cluster network as well as for end users who want to use that cluster.

This guide assumes that all users are familiar with cluster computing, that the cluster administrator is familiar with Linux® administration, and that the application programmer is familiar with MPI, vFabrics, SRP, and Distributed SA.

Related Materials

- *QLogic InfiniBand®* Adapter Hardware Installation Guide
- *QLogic InfiniBand®* Fabric Software Installation Guide
- Release Notes

Documentation Conventions

This guide uses the following documentation conventions:

- **NOTE:** provides additional information.
- ***CAUTION!*** indicates the presence of a hazard that has the potential of causing damage to data or equipment.
- ***WARNING!!*** indicates the presence of a hazard that has the potential of causing personal injury.

- Text in blue font indicates a hyperlink (jump) to a figure, table, or section in this guide, and links to Web sites are shown in underlined blue. For example:
  - Table 9-2 lists problems related to the user interface and remote agent.
  - See "Installation Checklist" on page 3-6.
  - For more information, visit www.qlogic.com.
- Text in **bold** font indicates user interface elements such as a menu items, buttons, check boxes, or column headings. For example:
  - Click the **Start** button, point to **Programs**, point to **Accessories**, and then click **Command Prompt**.
  - Under **Notification Options**, select the **Warning Alarms** check box.
- Text in `Courier` font indicates a file name, directory path, or command line text. For example:
  - To return to the root directory from anywhere in the file structure: Type `cd /root` and press ENTER.
  - Enter the following command: `sh ./install.bin`
- Key names and key strokes are indicated with UPPERCASE:
  - Press CTRL+P.
  - Press the UP ARROW key.
- Text in *italics* indicates terms, emphasis, variables, or document titles. For example:
  - For a complete listing of license agreements, refer to the *QLogic Software End User License Agreement*.
  - What are *shortcut keys*?
  - To enter the date type *mm/dd/yyyy* (where *mm* is the month, *dd* is the day, and *yyyy* is the year).
- Topic titles between quotation marks identify related topics either within this manual or in the online help, which is also referred to as *the help system* throughout this document.

License Agreements

Refer to the *QLogic Software End User License Agreement* for a complete listing of all license agreements affecting this product.

Technical Support

Customers should contact their authorized maintenance provider for technical support of their QLogic products. QLogic-direct customers may contact QLogic Technical Support; others will be redirected to their authorized maintenance provider. Visit the QLogic support Web site listed in Contact Information for the latest firmware and software updates.

For details about available service plans, or for information about renewing and extending your service, visit the Service Program web page at http://www.qlogic.com/services.

Training

QLogic offers training for technical professionals for all iSCSI, InfiniBand® (IB), and Fibre Channel products. From the main QLogic web page at www.qlogic.com, click the **Support** tab at the top, and then click **Training and Certification** on the left. The QLogic Global Training portal offers online courses, certification exams, and scheduling of in-person training.

Technical Certification courses include installation, maintenance and troubleshooting QLogic products. Upon demonstrating knowledge using live equipment, QLogic awards a certificate identifying the student as a certified professional. You can reach the training professionals at QLogic by e-mail at training@qlogic.com.

Contact Information

QLogic Technical Support for products under warranty is available during local standard working hours excluding QLogic Observed Holidays. For customers with extended service, consult your plan for available hours.For Support phone numbers, see the Contact Support link at support@qlogic.com.

Support Headquarters	QLogic Corporation 4601 Dean Lakes Blvd. Shakopee, MN 55379 USA
QLogic Web Site	www.qlogic.com
Technical Support Web Site	http://support.qlogic.com
Technical Support E-mail	support@qlogic.com
Technical Training E-mail	training@qlogic.com

Knowledge Database

The QLogic knowledge database is an extensive collection of QLogic product information that you can search for specific solutions. We are constantly adding to the collection of information in our database to provide answers to your most urgent questions. Access the database from the QLogic Support Center: http://support.qlogic.com.

1 Introduction

How this Guide is Organized

The *QLogic OFED+ Host Software User Guide* is organized into these sections:

- Section 1, provides an overview and describes interoperability.
- Section 2, describes how to setup your cluster to run high-performance MPI jobs.
- Section 3, describes the lower levels of the supplied QLogic OFED+ Host software. This section is of interest to a InfiniBand® cluster administrator.
- Section 4, helps the Message Passing Interface (MPI) programmer make the best use of the Open MPI implementation. Examples are provided for compiling and running MPI programs.
- Section 5, gives examples for compiling and running MPI programs with other MPI implementations.
- Section 7, describes QLogic Performance Scaled Messaging (PSM) that provides support for full Virtual Fabric (vFabric) integration, allowing users to specify InfiniBand® Service Level (SL) and Partition Key (PKey), or to provide a configured Service ID (SID) to target a vFabric.
- Section 8, describes dispersive routing in the InfiniBand® fabric to avoid congestion hotspots by “sraying” messages across the multiple potential paths.
- Section 9, describes open-source Preboot Execution Environment (gPXE) boot including installation and setup.
- Appendix A, describes how to run QLogic’s performance measurement programs.
- Appendix B, describes SCSI RDMA Protocol (SRP) configuration that allows the SCSI protocol to run over InfiniBand® for Storage Area Network (SAN) usage.

- Appendix C, describes two methods the administrator can use to allow users to submit MPI jobs through batch queuing systems.
- Appendix D, provides information for troubleshooting installation, cluster administration, and MPI.
- Appendix E, provides information for troubleshooting the upper layer protocol utilities in the fabric.
- Appendix F, provides instructions for checking write combining and for using the Page Attribute Table (PAT) and Memory Type Range Registers (MTRR).
- Appendix G, contains useful programs and files for debugging, as well as commands for common tasks.
- Appendix H, contains a list of useful web sites and documents for a further understanding of the InfiniBand® fabric, and related information.

In addition, the *QLogic InfiniBand®* Adapter Hardware Installation Guide contains information on QLogic hardware installation and the *QLogic InfiniBand®* Fabric Software Installation Guide contains information on QLogic software installation.

Overview

The material in this documentation pertains to a QLogic OFED+ cluster. A cluster is defined as a collection of nodes, each attached to an InfiniBand®-based fabric through the QLogic interconnect.

The QLogic IB Host Channel Adapters (HCA) are InfiniBand® 4X adapters. The quad data rate (QDR) adapters (QLE7340, QLE7342, QMH7342, and QME7342) have a raw data rate of 40Gbps (data rate of 32Gbps). The QLE7340, QLE7342, QMH7342, and QME7342 adapters can also run in DDR or SDR mode.

The QLogic IB HCA utilize standard, off-the-shelf InfiniBand® 4X switches and cabling. The QLogic interconnect is designed to work with all InfiniBand®-compliant switches.

NOTE

If you are using the QLE7300 series adapters in QDR mode, a QDR switch must be used.

QLogic OFED+ software is interoperable with other vendors' IBTA compliant InfiniBand® adapters running compatible OFED releases. There are several options for subnet management in your cluster:

- An embedded subnet manager can be used in one or more managed switches. QLogic offers the QLogic Embedded Fabric Manager (FM) for both DDR and QDR switch product lines supplied by your IB switch vendor.
- A host-based subnet manager can be used. QLogic provides the QLogic Fabric Manager (FM), as a part of the QLogic InfiniBand® Fabric Suite (IFS).

Interoperability

QLogic OFED+ participates in the standard IB subnet management protocols for configuration and monitoring. Note that:

- QLogic OFED+ (including Internet Protocol over InfiniBand® (IPoIB)) is interoperable with other vendors' InfiniBand® adapters running compatible OFED releases.
- In addition to supporting running MPI over verbs, QLogic provides a high-performance InfiniBand®-Compliant vendor-specific protocol, known as PSM. MPIs run over PSM will not interoperate with other adapters.

NOTE

See the OpenFabrics web site at www.openfabrics.org for more information on the OpenFabrics Alliance.

2 Step-by-Step Cluster Setup and MPI Usage Checklists

This section describes how to set up your cluster to run high-performance Message Passing Interface (MPI) jobs.

Cluster Setup

Perform the following tasks when setting up the cluster. These include BIOS, adapter, and system settings.

1. Make sure that hardware installation has been completed according to the instructions in the *QLogic InfiniBand*® Adapter Hardware Installation Guide and software installation and driver configuration has been completed according to the instructions in the *QLogic InfiniBand*® Fabric Software Installation Guide. To minimize management problems, the compute nodes of the cluster must have very similar hardware configurations and identical software installations. See “Homogeneous Nodes” on page 3-37 for more information.

2. Check that the BIOS is set properly according to the instructions in the *QLogic InfiniBand*® Adapter Hardware Installation Guide.

3. Set up the Distributed Subnet Administration (SA) to correctly synchronize your virtual fabrics. See “QLogic Distributed Subnet Administration” on page 3-12

4. Adjust settings, including setting the appropriate MTU size. See “Adapter and Other Settings” on page 3-38.

5. Remove unneeded services. See “Remove Unneeded Services” on page 3-39.

6. Disable powersaving features. See “Host Environment Setup for MPI” on page 3-40.

7. Check other performance tuning settings. See “Performance Settings and Management Tips” on page 3-24.

8. Set up the host environment to use `ssh`. Two methods are discussed in “Host Environment Setup for MPI” on page 3-40.

9. Verify the cluster setup. See “Checking Cluster and Software Status” on page 3-44.

Using MPI

1. Verify that the QLogic hardware and software has been installed on all the nodes you will be using, and that `ssh` is set up on your cluster (see all the steps in the Cluster Setup checklist).

2. Setup Open MPI. See “Setup” on page 4-2.

3. Compile Open MPI applications. See “Compiling Open MPI Applications” on page 4-2

4. Create an `mpihosts` file that lists the nodes where your programs will run. See “Create the mpihosts File” on page 4-3.

5. Run Open MPI applications. See “Running Open MPI Applications” on page 4-3.

6. Configure MPI programs for Open MPI. See “Configuring MPI Programs for Open MPI” on page 4-5

7. To test using other MPIs that run over PSM, such as MVAPICH, MVAPICH2, Platform MPI, and Intel MPI, see Section 5 Using Other MPIs.

8. To switch between multiple versions of MVAPICH, use the `mpi-selector`. See “Managing MVAPICH, and MVAPICH2 with the mpi-selector Utility” on page 5-5.

9. Refer to “Performance Tuning” on page 3-25 to read more about runtime performance tuning.

10. Refer to Section 5 Using Other MPIs to learn about using other MPI implementations.

3 InfiniBand® Cluster Setup and Administration

This section describes what the cluster administrator needs to know about the QLogic OFED+ software and system administration.

Introduction

The IB driver ib_qib, QLogic Performance Scaled Messaging (PSM), accelerated Message-Passing Interface (MPI) stack, the protocol and MPI support libraries, and other modules are components of the QLogic OFED+ software. This software provides the foundation that supports the MPI implementation.

Figure 3-1 illustrates these relationships. Note that HP-MPI, Platform MPI, Intel MPI, MVAPICH, MVAPICH2, and Open MPI can run either over PSM or OpenFabrics® User Verbs.

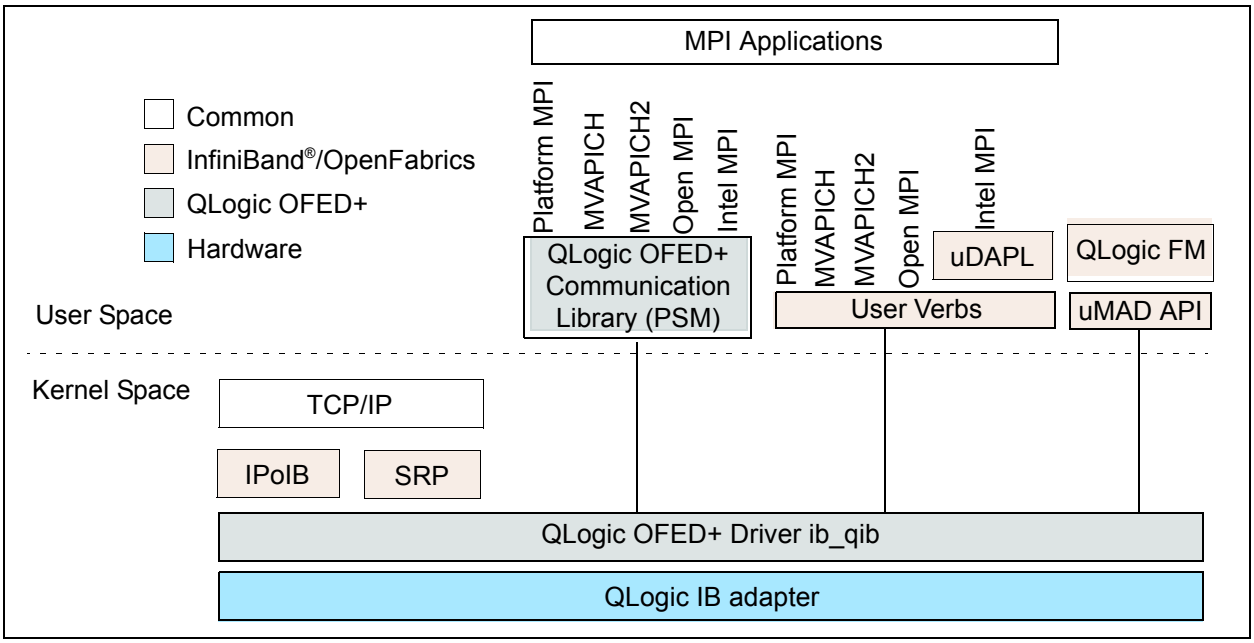

Figure 3-1. QLogic OFED+ Software Structure

Installed Layout

This section describes the default installed layout for the QLogic OFED+ software and QLogic-supplied MPIs.

QLogic-supplied Open MPI, MVAPICH, and MVAPICH2 RPMs with PSM support and compiled with GCC, PGI, and the Intel compilers are installed in directories using the following format:

`/usr/mpi/<compiler>/<mpi>-<mpi_version>-qlc`

For example:

`/usr/mpi/gcc/openmpi-1.4-qlc`

QLogic OFED+ utility programs, are installed in:

`/usr/bin`
`/sbin`
`/opt/iba/*`

Documentation is found in:

`/usr/share/man`
`/usr/share/doc/infinipath`

License information is found only in `usr/share/doc/infinipath`. QLogic OFED+ Host Software user documentation can be found on the QLogic web site on the software download page for your distribution.

Configuration files are found in:

`/etc/sysconfig`

Init scripts are found in:

`/etc/init.d`

The IB driver modules in this release are installed in:

```
/lib/modules/$(uname -r)/
updates/kernel/drivers/infiniband/hw/qib
```

Most of the other OFED modules are installed under the `infiniband` subdirectory. Other modules are installed under:

`/lib/modules/$(uname -r)/updates/kernel/drivers/net`

The RDS modules are installed under:

`/lib/modules/$(uname -r)/updates/kernel/net/rds`

IB and OpenFabrics Driver Overview

The `ib_qib` module provides low-level QLogic hardware support, and is the base driver for both MPI/PSM programs and general OpenFabrics protocols such as IPoIB and sockets direct protocol (SDP). The driver also supplies the Subnet Management Agent (SMA) component.

The following is a list of the optional configurable OpenFabrics components and their default settings:

- IPoIB network interface. This component is required for TCP/IP networking for running IP traffic over the IB link. It is not running until it is configured.
- OpenSM. This component is disabled at startup. QLogic recommends using the QLogic Fabric Manager (FM), which is included with the IFS or optionally available within the QLogic switches. QLogic FM or OpenSM can be installed on one or more nodes with only one node being the master SM.
- SRP (OFED and QLogic modules). SRP is not running until the module is loaded and the SRP devices on the fabric have been discovered.
- MPI over uDAPL (can be used by Intel MPI). IPoIB must be configured before MPI over uDAPL can be set up.

Other optional drivers can now be configured and enabled, as described in “IPoIB Network Interface Configuration” on page 3-3.

Complete information about starting, stopping, and restarting the QLogic OFED+ services are in “Managing the ib_qib Driver” on page 3-21.

IPoIB Network Interface Configuration

The following instructions show you how to manually configure your OpenFabrics IPoIB network interface. QLogic recommends using the QLogic OFED+ Host Software Installation package or the iba_config tool. For larger clusters, FastFabric can be used to automate installation and configuration of many nodes. These tools automate the configuration of the IPoIB network interface. This example assumes that you are using `sh` or `bash` as your shell, all required QLogic OFED+ and OpenFabric’s RPMs are installed, and your startup scripts have been run (either manually or at system boot).

For this example, the IPoIB network is 10.1.17.0 (one of the networks reserved for private use, and thus not routable on the Internet), with a /8 host portion. In this case, the netmask must be specified.

This example assumes that no hosts files exist, the host being configured has the IP address 10.1.17.3, and DHCP is not used.

NOTE

Instructions are only for this static IP address case. Configuration methods for using DHCP will be supplied in a later release.

1. Type the following command (as a root user):

```
ifconfig ib0 10.1.17.3 netmask 0xffffff00
```

2. To verify the configuration, type:

```
ifconfig ib0
ifconfig ib1
```

The output from this command will be similar to:

```
ib0   Link encap:InfiniBand HWaddr
00:00:00:02:FE:80:00:00:00:00:00:00:00:00:00:00:00:00:00:
00
inet addr:10.1.17.3  Bcast:10.1.17.255 Mask:255.255.255.0
UP BROADCAST RUNNING MULTICAST  MTU:4096  Metric:1
RX packets:0 errors:0 dropped:0 overruns:0 frame:0
TX packets:0 errors:0 dropped:0 overruns:0 carrier:0
collisions:0 txqueuelen:128
RX bytes:0 (0.0 b)  TX bytes:0 (0.0 b)
```

3. Type:

```
ping -c 2 -b 10.1.17.255
```

The output of the `ping` command will be similar to the following, with a line for each host already configured and connected:

```
WARNING: pinging broadcast address
PING 10.1.17.255 (10.1.17.255) 517(84) bytes of data.
174 bytes from 10.1.17.3: icmp_seq=0 ttl=174 time=0.022
ms
64 bytes from 10.1.17.1: icmp_seq=0 ttl=64 time=0.070 ms
(DUP!)
64 bytes from 10.1.17.7: icmp_seq=0 ttl=64 time=0.073 ms
(DUP!)
```

The IPoIB network interface is now configured.

4. Restart (as a root user) by typing:

```
/etc/init.d/openibd restart
```

NOTE

- The configuration must be repeated each time the system is rebooted.
- IPoIB-CM (Connected Mode) is enabled by default. The setting in `/etc/infiniband/openib.conf` is `SET_IPOIB_CM=yes`. To use datagram mode, change the setting to `SET_IPOIB_CM=no`. Setting can also be changed when asked during initial installation (./INSTALL).

IPoIB Administration

Administering IPoIB

Stopping, Starting and Restarting the IPoIB Driver

QLogic recommends using the QLogic IFS Installer TUI or iba_config command to enable autostart for the IPoIB driver. Refer to the *QLogic InfiniBand*® Fabric Software Installation Guide for more information. For using the command line to stop, start, and restart the IPoIB driver use the following commands.

To stop the IPoIB driver, use the following command:

```
/etc/init.d/openibd stop
```

To start the IPoIB driver, use the following command:

```
/etc/init.d/openibd start
```

To restart the IPoIB driver, use the following command:

```
/etc/init.d/openibd restart
```

Configuring IPoIB

QLogic recommends using the QLogic IFS Installer TUI, FastFabric, or `iba_config` command to configure the boot time and autostart of the IPoIB driver. Refer to the *QLogic InfiniBand*® Fabric Software Installation Guide for more information on using the QLogic IFS Installer TUI. Refer to the *QLogic FastFabric User Guide* for more information on using FastFabric. For using the command line to configure the IPoIB driver use the following commands.

Editing the IPoIB Configuration File

1. For each IP Link Layer interface, create an interface configuration file, `/etc/sysconfig/network/ifcfg-`*`NAME`*, where *`NAME`* is the value of the

NAME field specified in the `CREATE` block. The following is an example of the `ifcfg-`*NAME* file:

```
DEVICE=ib1
BOOTPROTO=static
BROADCAST=192.168.18.255
IPADDR=192.168.18.120
NETMASK=255.255.255.0
ONBOOT=yes
NM_CONTROLLED=no
```

NOTE

For IPoIB, the **`INSTALL`** script for the adapter now helps the user create the `ifcfg` files.

2. After modifying the `/etc/sysconfig/ipoib.cfg` file, restart the IPoIB driver with the following:

```
/etc/init.d/openibd restart
```

IB Bonding

IB bonding is a high availability solution for IPoIB interfaces. It is based on the Linux Ethernet Bonding Driver and was adopted to work with IPoIB. The support for IPoIB interfaces is only for the active-backup mode, other modes should not be used. QLogic supports bonding across HCA ports and bonding port 1 and port 2 on the same HCA.

Interface Configuration Scripts

Create interface configuration scripts for the ib*X* and bond*X* interfaces. Once the configurations are in place, perform a server reboot, or a service network restart. For SLES operating systems (OS), a server reboot is required. Refer to the following standard syntax for bonding configuration by the OS.

NOTE

For all of the following OS configuration script examples that set MTU, `MTU=65520` is valid only if all IPoIB slaves operate in connected mode and are configured with the same value. For IPoIB slaves that work in datagram mode, use MTU=2044. If the MTU is not set correctly or the MTU is not set at all (set to the default value), performance of the interface may be lower.

Red Hat EL5 and EL6

The following is an example for `bond0` (master). The file is named `/etc/sysconfig/network-scripts/ifcfg-bond0`:

```
DEVICE=bond0
IPADDR=192.168.1.1
NETMASK=255.255.255.0
NETWORK=192.168.1.0
BROADCAST=192.168.1.255
ONBOOT=yes
BOOTPROTO=none
USERCTL=no
MTU=65520
BONDING_OPTS="primary=ib0 updelay=0 downdelay=0"
```

The following is an example for `ib0` (slave). The file is named `/etc/sysconfig/network-scripts/ifcfg-ib0`:

```
DEVICE=ib0
USERCTL=no
ONBOOT=yes
MASTER=bond0
SLAVE=yes
BOOTPROTO=none
TYPE=InfiniBand
PRIMARY=yes
```

The following is an example for `ib1` (slave 2). The file is named `/etc/sysconfig/network-scripts/ifcfg-ib1`:

```
DEVICE=ib1
USERCTL=no
ONBOOT=yes
MASTER=bond0
SLAVE=yes
BOOTPROTO=none
TYPE=InfiniBand
```

Add the following lines to the RHEL 5.x file `/etc/modprobe.conf`, or the RHEL 6.x file `/etc/modprobe.d/ib_qib.conf`:

```
alias bond0 bonding
options bond0 miimon=100 mode=1 max_bonds=1
```

SuSE Linux Enterprise Server (SLES) 10 and 11

The following is an example for `bond0` (master). The file is named `/etc/sysconfig/network-scripts/ifcfg-bond0`:

```
DEVICE="bond0"
TYPE="Bonding"
IPADDR="192.168.1.1"
NETMASK="255.255.255.0"
NETWORK="192.168.1.0"
BROADCAST="192.168.1.255"
BOOTPROTO="static"
USERCTL="no"
STARTMODE="onboot"
BONDING_MASTER="yes"
BONDING_MODULE_OPTS="mode=active-backup miimon=100
primary=ib0 updelay=0 downdelay=0"
BONDING_SLAVE0=ib0
BONDING_SLAVE1=ib1
MTU=65520
```

The following is an example for `ib0` (slave). The file is named `/etc/sysconfig/network-scripts/ifcfg-ib0`:

```
DEVICE='ib0'
BOOTPROTO='none'
STARTMODE='off'
WIRELESS='no'
ETHTOOL_OPTIONS=''
NAME=''
USERCONTROL='no'
IPOIB_MODE='connected'
```

The following is an example for `ib1` (slave 2). The file is named `/etc/sysconfig/network-scripts/ifcfg-ib1`:

```
DEVICE='ib1'
BOOTPROTO='none'
STARTMODE='off'
WIRELESS='no'
ETHTOOL_OPTIONS=''
NAME=''
USERCONTROL='no'
IPOIB_MODE='connected'
```

Verify the following line is set to the value of yes in `/etc/sysconfig/boot`:

```
RUN_PARALLEL="yes"
```

Verify IB Bonding is Configured

After the configuration scripts are updated, and the service network is restarted or a server reboot is accomplished, use the following CLI commands to verify that IB bonding is configured.

- `cat /proc/net/bonding/bond0`
- `# ifconfig`

Example of `cat /proc/net/bonding/bond0` output:

```
# cat /proc/net/bonding/bond0
Ethernet Channel Bonding Driver: v3.2.3 (December 6, 2007)

Bonding Mode: fault-tolerance (active-backup) (fail_over_mac)
Primary Slave: ib0
Currently Active Slave: ib0
MII Status: up
MII Polling Interval (ms): 100
Up Delay (ms): 0
Down Delay (ms): 0

Slave Interface: ib0
MII Status: up
Link Failure Count: 0
Permanent HW addr: 80:00:04:04:fe:80

Slave Interface: ib1
MII Status: up
Link Failure Count: 0
Permanent HW addr: 80:00:04:05:fe:80
```

Example of `ifconfig` output:

```
st2169:/etc/sysconfig # ifconfig
bond0     Link encap:InfiniBand  HWaddr
80:00:00:02:FE:80:00:00:00:00:00:00:00:00:00:00:00:00:00:00
       inet addr:192.168.1.1  Bcast:192.168.1.255
Mask:255.255.255.0
       inet6 addr: fe80::211:7500:ff:909b/64 Scope:Link
       UP BROADCAST RUNNING MASTER MULTICAST  MTU:65520  Metric:1
       RX packets:120619276 errors:0 dropped:0 overruns:0 frame:0
       TX packets:120619277 errors:0 dropped:137 overruns:0
carrier:0
       collisions:0 txqueuelen:0
       RX bytes:10132014352 (9662.6 Mb)  TX bytes:10614493096
(10122.7 Mb)

ib0     Link encap:InfiniBand  HWaddr
80:00:00:02:FE:80:00:00:00:00:00:00:00:00:00:00:00:00:00:00
       UP BROADCAST RUNNING SLAVE MULTICAST  MTU:65520  Metric:1
       RX packets:118938033 errors:0 dropped:0 overruns:0 frame:0
       TX packets:118938027 errors:0 dropped:41 overruns:0
carrier:0
       collisions:0 txqueuelen:256
       RX bytes:9990790704 (9527.9 Mb)  TX bytes:10466543096
(9981.6 Mb)

ib1     Link encap:InfiniBand  HWaddr
80:00:00:02:FE:80:00:00:00:00:00:00:00:00:00:00:00:00:00:00
       UP BROADCAST RUNNING SLAVE MULTICAST  MTU:65520  Metric:1
       RX packets:1681243 errors:0 dropped:0 overruns:0 frame:0
       TX packets:1681250 errors:0 dropped:96 overruns:0 carrier:0
       collisions:0 txqueuelen:256
      RX bytes:141223648 (134.6 Mb)  TX bytes:147950000 (141.0 Mb)
```

Subnet Manager Configuration

QLogic recommends using the QLogic Fabric Manager to manage your fabric. Refer to the *QLogic Fabric Manager User Guide* for information on configuring the QLogic Fabric Manager.

OpenSM is a component of the OpenFabrics project that provides a Subnet Manager (SM) for IB networks. This package can optionally be installed on any machine, but only needs to be enabled on the machine in the cluster that will act as a subnet manager. You cannot use OpenSM if any of your IB switches provide a subnet manager, or if you are running a host-based SM, for example the QLogic Fabric Manager.

⚠WARNING

Don’t run OpenSM with QLogic FM in the same fabric.

If you are using the Installer tool, you can set the OpenSM default behavior at the time of installation.

OpenSM only needs to be enabled on the node that acts as the subnet manager.Toenable OpenSM the `iba_config` command can be used or the `chkconfig` command (as a root user) can be used on the node where it will be run. The `chkconfig` command to enable the OpenSM is:

```
chkconfig opensmd on
```

The `chkconfig` command to disable it on reboot is:

```
chkconfig opensmd off
```

You can start `opensmd` without rebooting your machine by typing:

```
/etc/init.d/opensmd start
```

You can stop `opensmd b`y typing:

```
/etc/init.d/opensmd stop
```

If you want to pass any arguments to the OpenSM program, modify the following file, and add the arguments to the `OPTIONS` variable:

```
/etc/init.d/opensmd
```

For example:

```
Use the UPDN algorithm instead of the Min Hop algorithm.
OPTIONS="-R updn"
```

For more information on OpenSM, see the OpenSM man pages, or look on the OpenFabrics web site.

QLogic Distributed Subnet Administration

As InfiniBand® clusters are scaled into the Petaflop range and beyond, a more efficient method for handling queries to the Fabric Manager is required. One of the issues is that while the Fabric Manager can configure and operate that many nodes, under certain conditions it can become overloaded with queries from those same nodes.

For example, consider an IB fabric consisting of 1,000 nodes, each with 4 processors. When a large MPI job is started across the entire fabric, each process needs to collect IB path records for every other node in the fabric - and every single process is going to be querying the subnet manager for these path records at roughly the same time. This amounts to a total of 3.9 million path queries just to start the job.

In the past, MPI implementations have side-stepped this problem by hand crafting path records themselves, but this solution cannot be used if advanced fabric management techniques such as virtual fabrics and mesh/torus configurations are being used. In such cases, only the subnet manager itself has enough information to correctly build a path record between two nodes.

The Distributed Subnet Administration (SA) solves this problem by allowing each node to locally replicate the path records needed to reach the other nodes on the fabric. At boot time, each Distributed SA queries the subnet manager for information about the relevant parts of the fabric, backing off whenever the subnet manager indicates that it is busy. Once this information is in the Distributed SA's database, it is ready to answer local path queries from MPI or other IB applications. If the fabric changes (due to a switch failure or a node being added or removed from the fabric) the Distributed SA updates the affected portions of the database. The Distributed SA can be installed and run on any node in the fabric. It is only needed on nodes running MPI applications.

Applications that use Distributed SA

The QLogic PSM Library has been extended to take advantage of the Distributed SA. Therefore, all MPIs that use the QLogic PSM library can take advantage of the Distributed SA. Other applications must be modified specifically to take advantage of it. For developers writing applications that use the Distributed SA, refer to the header file `/usr/include/Infiniband/ofedplus_path.h` for information on using Distributed SA APIs. This file can be found on any node where the Distributed SA is installed. For further assistance please contact QLogic Support.

Virtual Fabrics and the Distributed SA

The IBTA standard states that applications can be identified by a Service ID (SID). The QLogic Fabric Manager uses SIDs to identify applications. One or more applications can be associated with a Virtual Fabric using the SID. The Distributed SA is designed to be aware of Virtual Fabrics, but to only store records for those Virtual Fabrics that match the SIDs in the Distributed SA's configuration file. The Distributed SA recognizes when multiple SIDs match the same Virtual Fabric and will only store one copy of each path record within a Virtual Fabric. SIDs that match more than one Virtual Fabric will be associated with a single Virtual Fabric. The Virtual Fabrics that do not match SIDs in the Distributed SA's database will be ignored.

Configuring the Distributed SA

In order to absolutely minimize the number of queries made by the Distributed SA, it is important to configure it correctly, both to match the configuration of the Fabric Manager and to exclude those portions of the fabric that will not be used by applications using the Distributed SA. The configuration file for the Distributed SA is named `/etc/sysconfig/iba/qlogic_sa.conf`.

Default Configuration

As shipped, the QLogic Fabric Manager creates a single virtual fabric, called “Default” and maps all nodes and Service IDs to it, and the Distributed SA ships with a configuration that lists a set of thirty-one SIDs, 0x1000117500000000 through 0x100011750000000f and 0x1 through 0xf. This results in an arrangement like the one shown in Figure 3-2

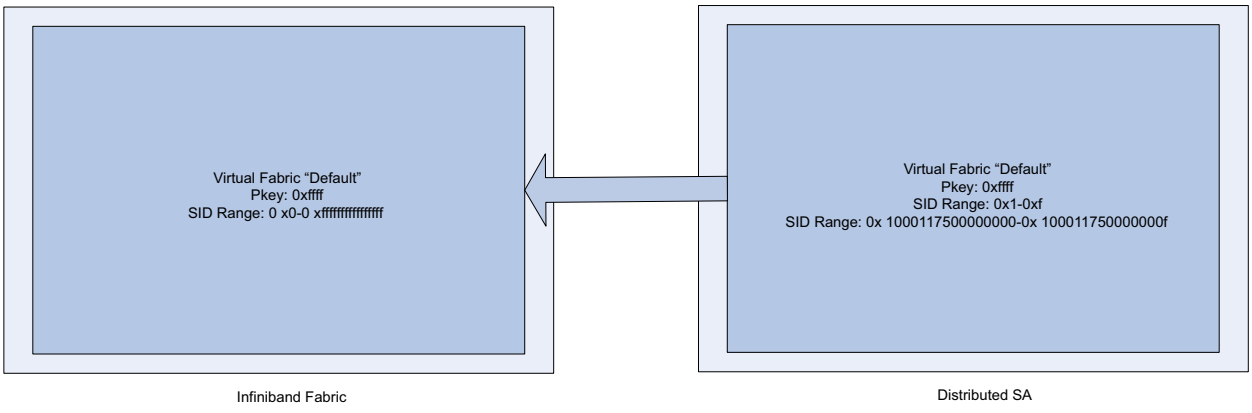

Figure 3-2. Distributed SA Default Configuration

If you are using the QLogic Fabric Manager in its default configuration, and you are using the standard QLogic PSM SIDs, this arrangement will work fine and you will not need to modify the Distributed SA's configuration file - but notice that the Distributed SA has restricted the range of SIDs it cares about to those that were defined in its configuration file. Attempts to get path records using other SIDs will not work, even if those other SIDs are valid for the fabric. When using this default configuration it is necessary that MPI applications only be run using one of these 32 SIDs.

Multiple Virtual Fabrics Example

A person configuring the physical IB fabric may want to limit how much IB bandwidth MPI applications are permitted to consume. In that case, they may re-configure the QLogic Fabric Manager, turning off the “Default” Virtual Fabric and replacing it with several other Virtual Fabrics.

In Figure 3-3, the administrator has divided the physical fabric into four virtual fabrics: “Admin” (used to communicate with the Fabric Manager), “Storage” (used by SRP), “PSM_MPI” (used by regular MPI jobs) and a special “Reserved” fabric for special high-priority jobs.

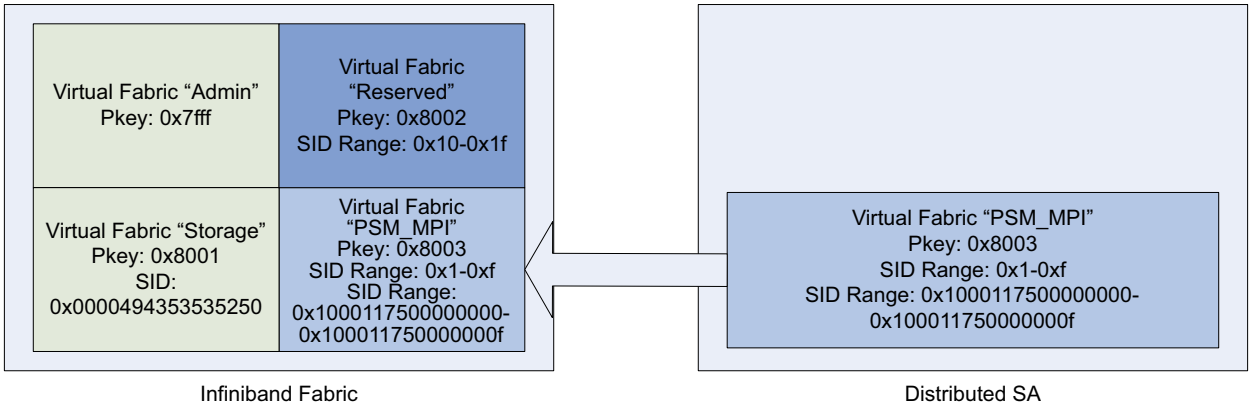

Figure 3-3. Distributed SA Multiple Virtual Fabrics Example

Due to the fact that the Distributed SA was not configured to include the SID Range 0x10 through 0x1f, it has simply ignored the “Reserved” VF. Adding those SIDs to the qlogic_sa.conf file solves the problem as shown in Figure 3-4.

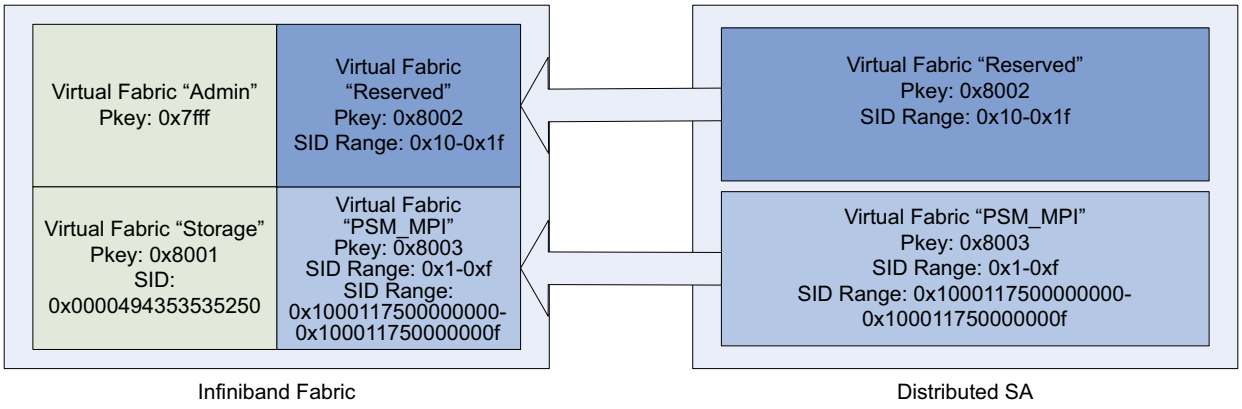

Figure 3-4. Distributed SA Multiple Virtual Fabrics Configured Example

Virtual Fabrics with Overlapping Definitions

As defined, SIDs should never be shared between Virtual Fabrics. Unfortunately, it is very easy to accidentally create such overlaps. Figure 3-5 shows an example with overlapping definitions.

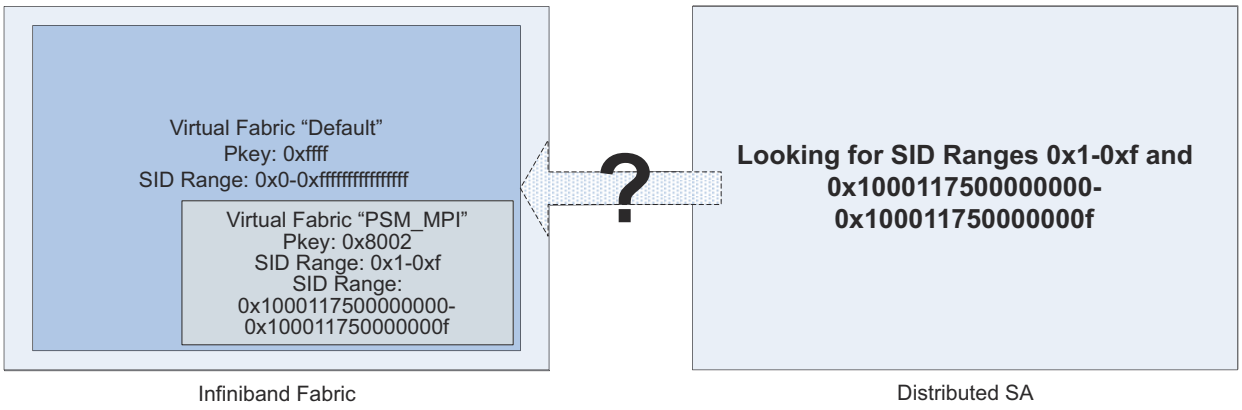

Figure 3-5. Virtual Fabrics with Overlapping Definitions

In Figure 3-5, the fabric administrator enabled the “PSM_MPI” Virtual Fabric without modifying the “Default” Virtual Fabric. As a result, the Distributed SA sees two different virtual fabrics that match its configuration file.

In Figure 3-6, the person administering the fabric has created two different Virtual Fabrics without turning off the Default - and two of the new fabrics have overlapping SID ranges.

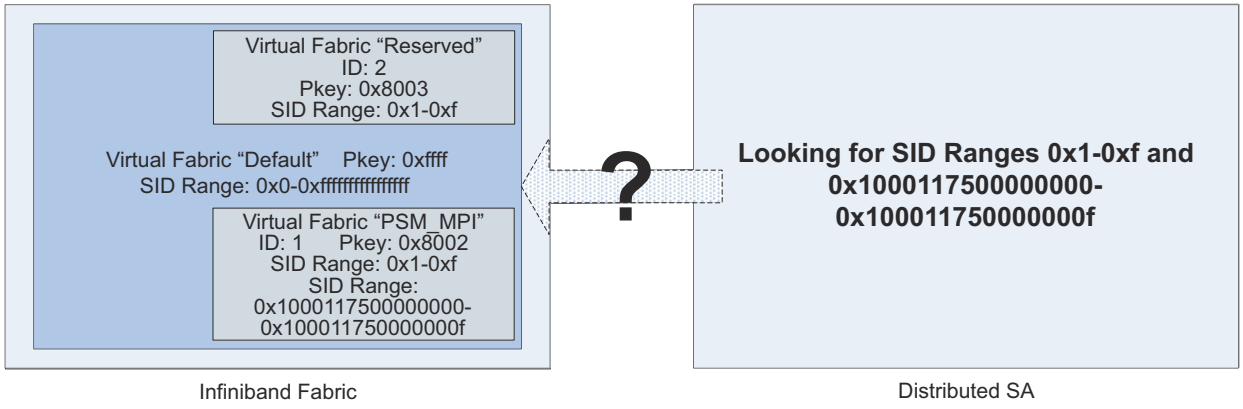

Figure 3-6. Virtual Fabrics with PSM_MPI Virtual Fabric Enabled

In Figure 3-6, the administrator enabled the “PSM_MPI” fabric, and then added a new “Reserved” fabric that uses one of the SID ranges that “PSM_MPI” uses. When a path query has been received, the Distributed SA deals with these conflicts as follows:

First, any virtual fabric with a pkey of 0xffff or 0x7fff is considered to be an Admin or Default virtual fabric. This Admin or Default virtual fabric is treated as a special case by the Distributed SA and is used only as a last resort. Stored SIDs are only mapped to the default virtual fabric if they do not match any other Virtual Fabrics. Thus, in the first example, Figure 3-6, the Distributed SA will assign all the SIDs in its configuration file to the “PSM_MPI” Virtual Fabric as shown in Figure 3-7.

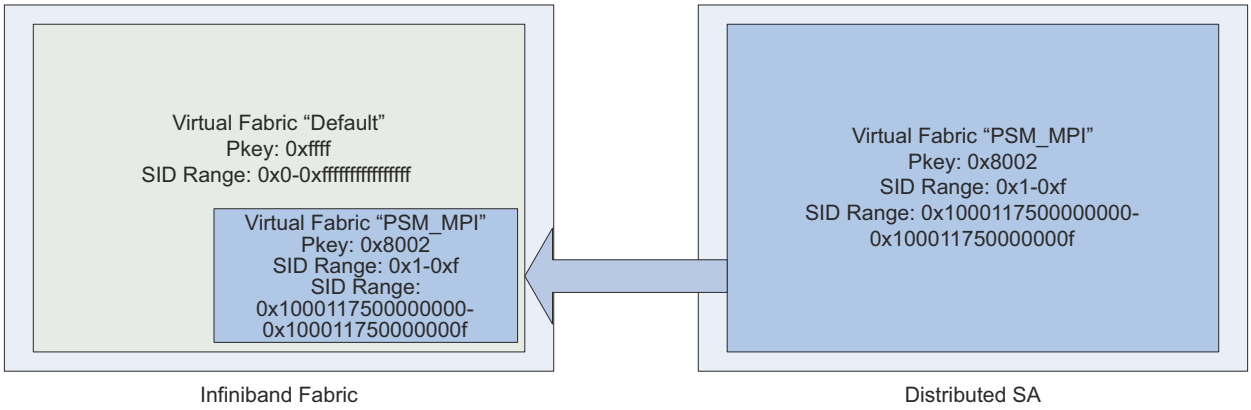

Figure 3-7. Virtual Fabrics with all SIDs assigned to PSM_MPI Virtual Fabric

Second, the Distributed SA handles overlaps by taking advantage of the fact that Virtual Fabrics have unique numeric indexes. These indexes are assigned by the QLogic Fabric Manager in the order which the Virtual Fabrics appear in the configuration file. These indexes can be seen by using the command `iba_saquery -o vfinfo` command. The Distributed SA will always assign a SID to the Virtual Fabric with the lowest index, as shown in Figure 3-8. This ensures that all copies of the Distributed SA in the IB fabric will make the same decisions about assigning SIDs. However, it also means that the behavior of your fabric can be affected by the order you configured the virtual fabrics.

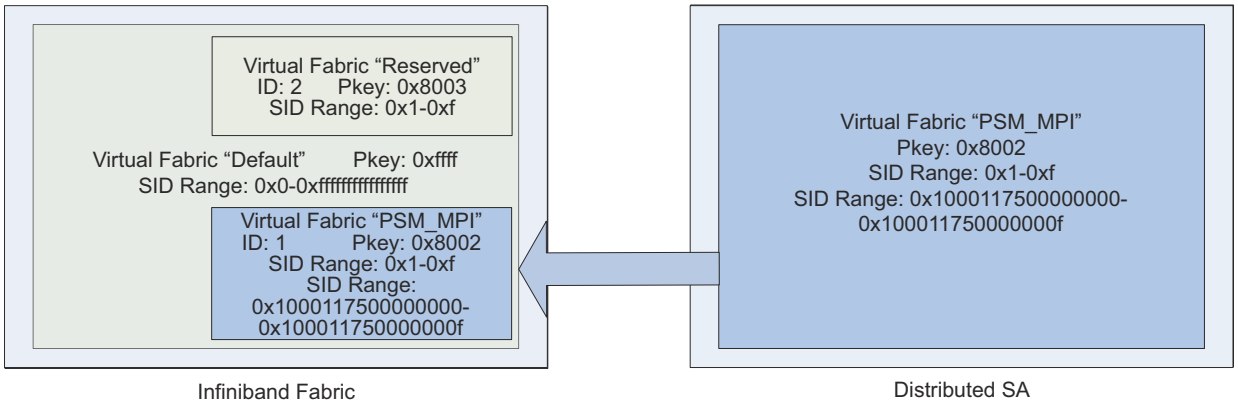

Figure 3-8. Virtual Fabrics with Unique Numeric Indexes

In Figure 3-8, the Distributed SA assigns all overlapping SIDs to the "PSM_MPI" fabric because it has the lowest Index

NOTE

The Distributed SA makes these assignments not because they are right, but because they allow the fabric to work even though there are configuration ambiguities. The correct solution in these cases is to redefine the fabric so that no node will ever be a member of two Virtual Fabrics that service the same SID.

Distributed SA Configuration File

The Distributed SA configuration file is `/etc/sysconfig/iba/qlogic_sa.conf`. It has several settings, but normally administrators will only need to deal with two or three of them.

SID

The SID is the primary configuration setting for the Distributed SA, and it can be specified multiple times. The SIDs identify applications which will use the distributed SA to determine their path records. The default configuration for the Distributed SA includes all the SIDs defined in the default Qlogic Fabric Manager configuration for use by MPI.

Each SID= entry defines one Service ID that will be used to identify an application. In addition, multiple SID= entries can be specified. For example, a virtual fabric has three sets of SIDs associated with it: 0x0a1 through 0x0a3, 0x1a1 through 0x1a3 and 0x2a1 through 0x2a3. You would define this as:

```
SID=0x0a1
SID=0x0a2
SID=0x0a3
SID=0x1a1
SID=0x1a2
SID=0x1a3
SID=0x2a1
SID=0x2a2
SID=0x2a3
```

NOTE

A SID of zero is not supported at this time. Instead, the OPP libraries treat zero values as "unspecified".

ScanFrequency

Periodically, the Distributed SA will completely re synchronize its database. This also occurs if the Fabric Manager is restarted. ScanFrequency defines the minimum number of seconds between complete re synchronizations. It defaults to 600 seconds, or 10 minutes. On very large fabrics, increasing this value can help reduce the total amount of SM traffic. For example, to set the interval to 15 minutes, add this line to the bottom of the qlogic_sa.conf file:

```
ScanFrequency=900
```

LogFile

Normally, the Distributed SA logs special events through syslog to `/var/log/messages`. This parameter allows you to specify a different destination for the log messages. For example, to direct Distributed SA messages to their own log, add this line to the bottom of the qlogic_sa.conf file:

```
LogFile=/var/log/SAReplica.log
```

Dbg

This parameter controls how much logging the Distributed SA will do. It can be set to a number between one and seven, where one indicates no logging and seven includes informational and debugging messages. To change the Dbg setting for Distributed SA, find the line in qlogic_sa.conf that reads `Dbg=5` and change it to a different value, between 1 and 7. The value of Dbg changes the amount of logging that the Distributed SA generates as follows:

- Dbg=1 or Dbg=2: Alerts and Critical Errors

  Only errors that will cause the Distributed SA to terminate will be reported.

- Dbg=3: Errors

  Errors will be reported, but nothing else. (Includes Dbg=1 and Dbg=2)

- Dbg=4: Warnings

  Errors and warnings will be reported. (Includes Dbg=3)

- Dbg=5: Normal

  Some normal events will be reported along with errors and warnings. (Includes Dbg=4)

- Dbg=6: Informational Messages

  In addition to the normal logging, Distributed SA will report detailed information about its status and operation. Generally, this will produce too much information for normal use. (Includes Dbg=5)

- Dbg=7: Debugging

  This should only be turned on at the request of QLogic Support. This will generate so much information that system operation will be impacted. (Includes Dbg=6)

Other Settings

The remaining configuration settings for the Distributed SA are generally only useful in special circumstances and are not needed in normal operation. The sample `qlogic_sa.conf` configuration file contains a brief description of each.

Changing the MTU Size

The Maximum Transfer Unit (MTU) size enabled by the IB HCA and set by the driver is 4KB. To see the current MTU size, and the maximum supported by the adapter, type the command:

```
$ ibv_devinfo
```

If the switches are set at 2K MTU size, then the HCA will automatically use this as the active MTU size, there is no need to change any file on the hosts.

To ensure that the driver on this host uses 2K MTU, add the following options line (as a root user) in to the configuration file:

```
options ib_qib ibmtu=4
```

Table 3-1 shows the value of each ibmtu number designation.

Table 3-1. ibmtu Values

Number Designation	Value in Bytes
1	256
2	512
3	1024
4	2048
5	4096

The following is a list of the configuration file locations for each OS:

- For RHEL 5.x use file: `/etc/modprobe.conf`
- For SLES 10 or 11 use file: `/etc/modprobe.conf.local`
- For RHEL 6.x use file: `/etc/modprobe.d/ib_qib.conf`

Restart the driver as described in Managing the ib_qib Driver.

NOTE

To use 4K MTU, set the switch to have the same 4K default. If you are using QLogic switches, the following applies:

- For the Externally Managed 9024, use 4.2.2.0.3 firmware (`9024DDR4KMTU_firmware.emfw`) for the 9024 EM. This has the 4K MTU default, for use on fabrics where 4K MTU is required. If 4K MTU support is not required, then use the `4.2.2.0.2 DDR *.emfw` file for DDR externally-managed switches. Use FastFabric (FF) to load the firmware on all the 9024s on the fabric.
- For the 9000 chassis, use the most recent 9000 code 4.2.4.0.1. The 4K MTU support is in 9000 chassis version 4.2.1.0.2 and later. For the 9000 chassis, when the FastFabric 4.3 (or later) chassis setup tool is used, the user is asked to select an MTU. FastFabric can then set that MTU in all the 9000 internally managed switches. The change will take effect on the next reboot. Alternatively, for the internally managed 9000s, the `ismChassisSetMtu` Command Line Interface (CLI) command can be used. This should be executed on every switch and both hemispheres of the 9240s.
- For the 12000 switches, refer to the *QLogic FastFabric User Guide* for externally managed switches, and to the *QLogic FastFabric CLI Reference Guide* for the internally managed switches.

For reference, see the *QLogic FastFabric User Guide* and the *QLogic 12000 CLI Reference Guide*. Both are available from the QLogic web site.

For other switches, see the vendors' documentation.

Managing the ib_qib Driver

The startup script for `ib_qib` is installed automatically as part of the software installation, and normally does not need to be changed. It runs as a system service.

The primary configuration file for the IB driver `ib_qib` and other modules and associated daemons is `/etc/infiniband/openib.conf`.

Normally, this configuration file is set up correctly at installation and the drivers are loaded automatically during system boot once the software has been installed. However, the `ib_qib` driver has several configuration variables that set reserved buffers for the software, define events to create trace records, and set the debug level.

If you are upgrading, your existing configuration files will not be overwritten.

See the `ib_qib` man page for more details.

Configure the ib_qib Driver State

Use the following commands to check or configure the state. These methods will not reboot the system.

To check the configuration state, use this command. You do not need to be a root user:

```
$ chkconfig --list openibd
```

To enable the driver, use the following command (as a root user):

```
# chkconfig openibd on 2345
```

To disable the driver on the next system boot, use the following command (as a root user):

```
# chkconfig openibd off
```

NOTE

This command does not stop and unload the driver if the driver is already loaded nor will it start the driver.

Start, Stop, or Restart ib_qib Driver

Restart the software if you install a new QLogic OFED+ Host Software release, change driver options, or do manual testing.

QLogic recommends using `/etc/init.d/openibd` to stop, stat and restart the ib_qib driver. For using the command line to stop, start, and restart (as a root user) the ib_qib driver use the following syntex:

```
# /etc/init.d/openibd [start | stop | restart]
```

⚠WARNING

If QLogic Fabric Manager, or OpenSM is configured and running on the node, it must be stopped before using the `openibd stop` command, and may be started after using the `openibd start` command.

This method will not reboot the system. The following set of commands shows how to use this script.

When you need to determine which ib_qib driver and OpenFabrics modules are running, use the following command. You do not need to be a root user.

```
$ lsmod | egrep 'ipath_|ib_|rdma_|findex'
```

You can check to see if `opensmd` is configured to autostart by using the following command (as a root user); if there is no output, `opensmd` is not configured to autostart:

```
# /sbin/chkconfig --list opensmd | grep -w on
```

Unload the Driver/Modules Manually

You can also unload the driver/modules manually without using `/etc/init.d/openibd`. Use the following series of commands (as a root user):

```
# umount /ipathfs
# fuser -k /dev/ipath* /dev/infiniband/*
# lsmod | egrep '^ib_|^rdma_|^iw_' | xargs modprobe -r
```

ib_qib Driver Filesystem

The ib_qib driver supplies a filesystem for exporting certain binary statistics to user applications. By default, this filesystem is mounted in the `/ipathfs` directory when the ib_qib script is invoked with the `start` option (e.g. at system startup). The filesystem is unmounted when the ib_qib script is invoked with the `stop` option (for example, at system shutdown).

Here is a sample layout of a system with two cards:

/ipathfs/0/flash

/ipathfs/0/port2counters

/ipathfs/0/port1counters

/ipathfs/0/portcounter_names

/ipathfs/0/counter_names

/ipathfs/0/counters

/ipathfs/driver_stats_names

/ipathfs/driver_stats

/ipathfs/1/flash

/ipathfs/1/port2counters

/ipathfs/1/port1counters

/ipathfs/1/portcounter_names

/ipathfs/1/counter_names

/ipathfs/1/counters

The `driver_stats` file contains general driver statistics. There is one numbered subdirectory per IB device on the system. Each numbered subdirectory contains the following per-device files:

- `port1counters`
- `port2counters`
- `flash`

The `driver1counters` and `driver2counters` files contain counters for the device, for example, interrupts received, bytes and packets in and out, etc. The `flash` file is an interface for internal diagnostic commands.

The file `counter_names` provides the names associated with each of the counters in the binary `port#counters` files, and the file `driver_stats_names` provides the names for the stats in the binary `driver_stats` files.

More Information on Configuring and Loading Drivers

See the `modprobe(8)`, `modprobe.conf(5)`, and `lsmod(8)` man pages for more information. Also see the file `/usr/share/doc/initscripts-*/sysconfig.txt` for more general information on configuration files.

Performance Settings and Management Tips

The following sections provide suggestions for improving performance and simplifying cluster management. Many of these settings will be done by the system administrator.

Performance Tuning

Tuning compute or storage (client or server) nodes with IB HCAs for MPI and verbs performance can be accomplished in several ways:

- Run the `ipath_perf_tuning` script in automatic mode (See "Performance Tuning using ipath_perf_tuning Tool" on page 3-34) (easiest method)
- Run the `ipath_perf_tuning` script in interactive mode (See "Performance Tuning using ipath_perf_tuning Tool" on page 3-34 or see `man ipath_perf_tuning`). This interactive mode allows more control, and should be used for tuning storage (client or server) nodes.
- Make changes to ib_qib driver parameter files, the BIOS or system services using the information provided in the following sections

NOTE

The modprobe configuration file (`modprobe.conf`) will be used in this section for the ib_qib module configuration file, which has various paths and names in the different Linux distributions as shown in the following list:

- For RHEL 5.x use file: `/etc/modprobe.conf`
- For SLES 10 or 11 use file: `/etc/modprobe.conf.local`
- For RHEL 6.x use file:`/etc/modprobe.d/ib_qib.conf`

Systems in General (With Either Intel or AMD CPUs)

For best performance on dual-port HCAs on which only one port is active, the module parameter line in the modprobe.conf file should include the following:

```
options ib_qib singleport=1
```

Services

Turn off the specified daemons using one of the following commands according to which OS is being used:

- For RHEL or similar systems use:

```
/sbin/chkconfig --level 12345 cpuspeed off
```

- For SLES systems use:

```
/sbin/chkconfig --level 12345 powersaved off
```

If `cpuspeed` or `powersaved` are being used as part of implementing Turbo modes to increase CPU speed, then they can be left on. With these daemons left on, IB micro-benchmark performance results may be more variable from run-to-run.

For compute nodes, set the default runlevel to 3 to reduce overheads due to unneeded processes. Reboot the system for this change to take effect.

Default Parameter Settings

The qib driver makes certain settings by default based on a check of which CPUs are in the system. Since these are done by default, no user- or ipath_perf_tuning-generated changes need to be made in the modprobe configuration file. It doesn't hurt anything if these settings are in the file, but they are not necessary.

On all systems, the qib driver behaves as if the following parameters were set:

```
rcvhdrcnt=4096
```

If you run a script, such as the following:

```
for x in /sys/module/ib_qib/parameters/*; do echo $(basename
$x) $(cat $x); done
```

Then in the list of qib parameters, you should see the following parameter being discussed:

```
. . .
rcvhdrcnt 0
```

The `0` means the driver automatically sets these parameters. Therefore, neither the user nor the `ipath_perf_tuning` script should modify these parameters.

Compute-only Node (Not part of a parallel file system cluster)

No tuning is required, other than what is in the Systems in General (With Either Intel or AMD CPUs) section.

For more details on settings that are specific to either Intel or AMD CPUs, refer to the following sections for details on systems with those types of CPUs.

Storage Node (for example, Lustre/GPFS client or server node)

Although termed a “Storage Node” this information includes nodes that are primarily compute nodes, but also act as clients of a parallel file server.

Increasing the number of kernel receive queues allows more CPU cores to be involved in the processing of verbs traffic. This is important when using parallel file systems such as Lustre or IBM's GPFS (General Parallel File System). The module parameter that sets this number is `krcvqs`. Each additional kernel receive queue (beyond the one default queue for each port) takes user contexts away from PSM and from the support of MPI or compute traffic. The formula which illustrates this trade-off is:

```
PSM Contexts = 16 - (krcvqs-1)x num_ports
```

Where *`number_ports`* is the number of ports on the HCA

For example, on a single-port card with `krcvqs=4` set in `modprobe.conf`:

```
PSM Contexts = 16 - (4-1)x 1 = 16 - 3 = 13
```

If this were a 12-core node, then 13 is more than enough PSM contexts to run an MPI process on each core without making use of context-sharing. An example , `ib_qib options` line in the `modprobe.conf` file, for this 12-core node case is:

```
options ib_qib singleport=1 krcvqs=4
```

Table 3-2 can be used as a guide for setting the `krcvqs` parameter for the number of cores in the system supporting PSM processes and the number of ports in the HCA. The table applies most readily to nodes with 1 HCA being used to support PSM (for example, MPI or SHMEM) processes. For nodes with multiple HCAs that are being used for PSM, the table decide the maximum number of cores that will be assigned on each HCA to support PSM (MPI or SHMEM) processes, then apply the table to each HCA in turn.

Table 3-2. `krcvqs` Parameter Settings

Cores per Node (to be used for MPI/PSM on 1 HCA):	1-port, Set krcvqs=	2 active ports in the HCA, Set krcvqs=
61-64	1[a]	1[a]
57-60	2	1[a]
53-56	3	2,1 (2 for port 1, 1 for one port)
12-52	4	2
8-11	3	2,1 (2 for port 1, 1 for one port)
4-7	2	1[a]
1-3	1[a]	1[a]

[a] 1 is the default setting, so if the table recommends '1', `krcvqs` does not need to be set.

In the rare case that the node has more than 64 cores, and it is desired to run MPI on more than 64 cores, then two HCAs are required and settings can be made, using the rules in Table 3-2, as though half the cores were assigned to each HCA.

AMD CPU Systems

To improve IPoIB and other verbs-based throughput performance, on AMD CPU systems, QLogic recommends setting pcie_caps=0x51 numa_aware=1 as modprobe configuration file parameters. For example, the module parameter line in the modprobe configuration file should include the following for AMD Opteron CPUs:

```
options ib_qib pcie_caps=0x51 numa_aware=1
```

On AMD systems, the pcie_caps=0x51 setting will result in a line of the `lspci -vv` output associated with the QLogic HCA reading in the "DevCtl" section:

MaxPayload 128 bytes, MaxReadReq 4096 bytes.

AMD Interlagos CPU Systems

With AMD Interlagos (Opteron 6200 Series) CPU systems, better performance will be obtained if, on single-HCA systems, the HCA is put in a PCIe slot closest to Socket number 1. You can typically find out which slots these are by looking at the schematics in the manual for your motherboard. (There is currently a BIOS or kernel problem which implies that no NUMA topology information is available from the kernel.)

To obtain top “Turbo boosts” of up to 1GHz in clock rate, when running on half the cores of a node, AMD recommends enabling the C6 C-state in the BIOS. Some applications (but certainly not all) run better when running on half the cores or a Interlagos node (on every other core, one per Bulldozer module). QLogic recommends enabling this C-state in the BIOS.

Intel CPU Systems

Typical tuning for recent Intel CPUs

For recent Intel CPUs (code-named Sandy Bridge, Westmere or Nehalem), set the following BIOS parameters:

- Disable all C-States.
- Disable Intel Hyper-Threading technology

For setting all C-States to 0 where there is no BIOS support:

1. Add kernel boot option using the following command:

```
processor.max_cstate=0
```

2. Reboot the system.

If the node uses a single-port HCA, and is not a part of a parallel file system cluster, there is no need for performance tuning changes to a modprobe configuration file. The driver will automatically set the parameters appropriately for the node's Intel CPU, in a conservative manner.

For all Intel systems with Xeon 5500 Series (Nehalem) or newer CPUs, the following settings are default:

```
pcie_caps=0x51
```

On Intel systems with Xeon 5500 Series (Nehalem) or newer CPUs, the `lspci` output will read:

```
MaxPayload 256 bytes, MaxReadReq 4096 bytes
```

If you run a script, such as the following:

```
for x in /sys/module/ib_qib/parameters/*; do echo $(basename
$x) $(cat $x); done
```

Then in the list of qib parameters, you should see the following for the two parameters being discussed:

```
. . .
rcvhdrcnt 0
. . .
pcie_caps 0
```

The `0` means the driver automatically sets these parameters. Therefore, neither the user nor the `ipath_perf_tuning` script should modify these parameters.

Intel Nehalem or Westmere CPU Systems (DIMM Configuration)

Compute node memory bandwidth is important for high-performance computing (HPC) application performance and for storage node performance. On Intel CPUs code named Nehalem or Westmere (Xeon 5500 series or 5600 series) it is important to have an equal number of dual in-line memory modules (DIMMs) on each of the three memory channels for each CPU. On the common dual CPU systems, you should use a multiple of six DIMMs for best performance.

High Risk Tuning for Intel Harpertown CPUs

For tuning the Harpertown generation of Intel Xeon CPUs that entails a higher risk factor, but includes a bandwidth benefit, the following can be applied:

For nodes with Intel Harpertown, Xeon 54xx CPUs, you can add `pcie_caps=0x51` and `pcie_coalesce=1` to the `modprobe.conf` file. For example:

```
options ib_qib pcie_caps=0x51 pcie_coalesce=1
```

If the following problem is reported by syslog, a typical diagnostic can be performed, which is described in the following paragraphs:

```
[PCIe Poisoned TLP][Send DMA memory read]
```

Another potential issue is that after starting `openibd`, messages such as the following appear on the console:

```
Message from syslogd@st2019 at Nov 14 16:55:02 ...
kernel:Uhhuh. NMI received for unknown reason 3d on CPU 0
```

After this happens, you may also see the following message in the syslog:

```
Mth dd hh:mm:ss st2019 kernel: ib_qib 0000:0a:00.0:
infinipath0:
Fatal Hardware Error, no longer usable, SN AIB1013A43727
```

These problems typically occur on the first run of an MPI program running over the PSM transport or immediately after the link becomes active. The adapter will be unusable after this situation until the system is rebooted. To resolve this issue try the following solutions in order:

- Remove `pcie_coalesce=1`
- Restart `openibd` and try the MPI program again
- Remove both `pcie_caps=0x51` and `pcie_coalesce=1` options from the `ib_qib` line in `modprobe.conf` file and reboot the system

NOTE

Removing both options will technically avoid the problem but can result in an unnecessary performance decrease. If the system has already failed with the above diagnostic it will need to be rebooted. Note that in modprobe.conf file all options for a particular kernel module must be on the same line and not on repeated options ib_qib lines.

Additional Driver Module Parameter Tunings Available

Setting driver module parameters on Per-unit or Per-port basis

The `ib_qib` driver allows the setting of different driver parameter values for the individual HCAs and ports. This allows the user to specify different values for each port on a HCA or different values for each HCA in the system. This feature is used when there is a need to tune one HCA or port for a particular type of traffic, and a different HCA or port for another type of traffic, for example, compute versus storage traffic.

Not all driver parameters support per-unit or per-port values. The driver parameters which can be used with the new syntax are listed below:

Per-unit parameters:

- `singleport` – Use only IB port 1; more per-port buffer space
- `cfgctxts` – Set max number of contexts to use
- `pcie_caps` – Max PCIe tuning: MaxPayload, MaxReadReq

Per-port parameters:

- `ibmtu` – Set max IB MTU
- `krcvqs` – number of kernel receive queues
- `num_vls` – Set number of Virtual Lanes to use

Specifying individual unit/port values is done by using a specific module parameter syntax:

param name=[*default*,][*unit*[:*port*]=*value*]

Where:

- *param name* is the driver module parameter name (listed above)
- *default* is the default value for that parameter. This value will be used for all remaining units/port which have not had individual values set. If no individual unit/port values have been specified, the default value will be used for all units/ports
- *unit* is the index of the HCA unit (as seen by the driver). This value is 0-based (index of first unit is '0').
- *port* is the port number on that HCA. This value is 1-based (number of first port is '1').

- *value* is the parameter value for the particular unit or port.

The fields in the square brackets are options; however, either a default or a per-unit/per-port value is required.

Example usage:

To set the default IB MTU to 1K for all ports on all units:

```
ibmtu=3
```

To set the IB MTU to 256-bytes for unit 0/port 1 and 4096-bytes for unit 0/port 2:

```
ibmtu=0:1=1,0:2=5
```

To set the default IB MTU to 2K for all ports but specify 4K for unit 0/port 1:

```
ibmtu=4,0:1=5
```

To singleport to OFF as default and turn it ON for unit 1:

```
singleport=0,1=1
```

To set number of configured contexts to 10 on unit 0 and 16 on unit 1:

```
cfgctxts=0=10,1=16
```

A user can identify HCAs and correlate them to system unit numbers by using the `-b` option (beacon mode option) to the `ipath_control` script. Issuing the following command (as root):

```
ipath_control -u unit -b on
```

Where:

unit is the system unit number will cause that HCA to start blinking the LEDs on the face of the board in an alternating pattern.

Once the board has been identified, the user can return the LEDs to normal mode of operation with the following command (as root):

```
ipath_control -u unit -b off
```

numa_aware

The Non-Uniform Memory Access (NUMA) awareness (`numa_aware`) module parameter enables driver memory allocations in the same memory domain or NUMA node of the HCA. This improves the overall system efficiency with CPUs on the same NUMA node having faster access times and higher bandwidths to memory.

The default is:

```
option ib_qib numa_aware=10
```

This command lets the driver automatically decide on the allocation behavior and disables this feature on platforms with AMD and Intel Westmere-or-earlier CPUs, while enabling it on newer Intel CPUs.

Tunable options:

```
option ib_qib numa_aware=0
```

This command disables the NUMA awareness when allocating memories within the driver. The memory allocation requests will be satisfied on the node's CPU that executes the request.

```
option ib_qib numa_aware=1
```

This command enables this feature with the driver allocating memory on the NUMA node closest to the HCA.

recv_queue_size, Tuning Related to NAKs

The Receiver Not Ready Negative Acknowledgement (RNR NAKs) can slow IPoIB down significantly. IB is fast enough to overrun IPoIB's receive queue before the post receives can occur.

The counter to look for on the sending side in this file is `RC RNR NAKs` as shown in the following example:

```
# cat /sys/class/infiniband/qib0/stats
Port 1:
RC timeouts 0
RC resends  0
RC QACKs    0
RC SEQ NAKs 0
RC RDMA seq 0
RC RNR NAKs 151 <---------
RC OTH NAKs 0
. . .
Ctx:npkts 0:170642806
```

Check the RC RNR NAKs before and after running the IPoIB test to see if that counter is increasing. If so, then increasing IPoIB's `recv_queue_size` to 512 in the `ib_ipoib.conf` file should eliminate RNR NAKs.

For example:

```
# cat /etc/modprobe.d/ib_ipoib.conf
alias ib0 ib_ipoib
alias ib1 ib_ipoib
options ib_ipoib recv_queue_size=512
```

Performance Tuning using `ipath_perf_tuning` Tool

The `ipath_perf_tuning` tool is intended to adjust parameters to the IB QIB driver to optimize the IB and application performance. The tool is designed to be run once per installation, however it can be re-run if changes to the configuration need to be made. Changes are made to the appropriate modprobe file depending on Linux distribution (see Affected Files).

The tool takes into account the type of the node being configured and can be run in one of two modes - automatic (the default) and interactive. In automatic mode, the tool will make the parameter adjustments without the need for any user input. Interactive mode will prompt the user for input on some of the settings and actions.

Table 3-3 list the checks the tool performs on the system on which it is run.

Table 3-3. Checks Preformed by `ipath_perf_tuning` Tool

Check Type	Description
`pcie_caps`	Adjust PCIe tuning for max payload and read request size. The result of this test depends on the CPU type of the node.
`singleport`	Determine whether to run the HCA in single port mode increasing the internal HCA resources for that port. This setting depends on the user's input and is only performed in interactive mode.
`krcvqs`	Determine the number of kernel receive context to allocate. Normally, the driver allocates one context per physical port. However, more kernel receive contexts can be allocated to improve Verbs performance.
`pcie_coalesce`	Enable PCIe coalescing. PCIe coalescing is only needed or enabled on some systems with Intel Harpertown CPUs.
`cache_bypass_copy`	Enable the use of Cache bypass copies. This option is enabled on AMD Interlagos (62xx) series processors.
`numa_aware`	Enable NUMA-aware memory allocations. This option is enabled on AMD CPUs only.

Table 3-3. Checks Preformed by `ipath_perf_tuning` Tool

Check Type	Description
`cstates`	Check whether (and which) C-States are enabled. C-States should be turned off for best performance.
`services`	Check whether certain system services (daemons) are enabled. These services should be turned off for best performance.

The values picked for the various checks and tests may depend on the type of node being configured. The tool is aware of two types of nodes—compute and storage nodes.

Compute Nodes

Compute nodes are nodes which should be optimized for faster computation and communication with other compute nodes.

Storage (Client or Server) Nodes

Storage nodes are nodes which serve as clients or servers in a parallel filesystem network. Storage nodes (especially clients) are typically performing computation and using MPI, in addition to sending and receiving storage network traffic. The objective is to improve IB verbs communications while maintaining good MPI performance.

OPTIONS

Table 3-4 list the options for the ***`ipath_perf_tuning`*** tool and describes each option.

Table 3-4. `ipath_perf_tuning` Tool Options

Option	Description
`-h`	Display a short multi-line help message
`-T test`	This option is used to limit the list of tests/check which the tool performs to only those specified by the option. Multiple tests can be specified as a comma-separated list.
`-I`	Run the tool in interactive mode. In this mode, the tool will prompt the user for input on certain tests.

AUTOMATIC vs. INTERACTIVE MODE

The tool performs different functions when running in automatic mode compared to running in the interactive mode. The differences include the node type selection, test execution, and applying the results of the executed tests.

Node Type Selection

The tool is capable of configuring compute nodes or storage nodes (see Compute Nodes and Storage (Client or Server) Nodes). When the tool is executed in interactive mode, it will query the user for the type of node. When the tool is running in automatic mode, it assumes that the node being configured is a compute node.

Test Execution

The main difference between the two test modes is that some of the tests are effectively skipped when the tool is in automatic mode. This is done, due to the fact, that these test do not provide a guaranteed universal performance gain and therefore, changing driver parameters associated with them requires user approval. Other tests, where the tool can make a safe determination, are performed in both modes without any user interaction. Table 3-5 list the test and describe the mode(s) for each.

Table 3-5. Test Execution Modes

Test	Mode
`pcie_caps`	Test is performed in both modes without any user interaction.
`singleport`	Test is only performed in interactive mode. The user is queried whether to enable singleport mode.
`krcvqs`	Test is performed in both modes without any user interaction.
`pci_coalesce`	Test is performed only in interactive mode. The user is queried whether to enable PCIe coalescing.
`cache_bypass_copy`	Test is performed in both modes without any user interaction.
`num_aware`	Test is performed in both modes without any user interaction.
`cstates`	Test is performed in both modes but the user is only notified of a potential issue if the tool is in interactive mode. In that case, the tool displays a warning and a suggestion on how to fix the issue.

Table 3-5. Test Execution Modes

Test	Mode
`services`	Test is performed in both modes but the user is notified of running services only if the tool is in interactive mode. In that case, the user is queried whether to turn the services off.

Applying the Results

Automatic mode versus interactive mode also has an effect when the tool is committing the changes to the system. Along with the necessary driver parameters, the script also writes a comment line in the appropriate file which serves as a marker. This marker contains the version of the script which is making the changes. If the version recorded matches the version of the script currently being run, the changes are only committed if the tool is in interactive mode. The assumption is that the script is being re-run by the user to make adjustments.

Affected Files

The following lists the distribution and the file that is modified by the ***`ipath_perf_tuning`*** tool:

- RHEL 6.0 and later – `/etc/modprobe.d/ib_qib.conf`
- RHEL prior to 6.0 – `/etc/modprobe.conf`
- SLES – `/etc/modprobe.conf.local`

Homogeneous Nodes

To minimize management problems, the compute nodes of the cluster should have very similar hardware configurations and identical software installations. A mismatch between the software versions can also cause problems. Old and new libraries must not be run within the same job. It may also be useful to distinguish between the IB-specific drivers and those that are associated with kernel.org, OpenFabrics, or are distribution-built. The most useful tools are:

- `ident` (see “ident” on page G-24)
- `ipathbug-helper` (see “ipath_checkout” on page G-25)
- `ipath_checkout` (see “ipath_checkout” on page G-25)
- `ipath_control` (see “ipath_control” on page G-27)
- `mpirun` (see “mpirun” on page G-31)

- `rpm` (see “rpm” on page G-32)
- `strings` (see “strings” on page G-32)

NOTE

Run these tools to gather information before reporting problems and requesting support.

Adapter and Other Settings

The following adapter and other settings can be adjusted for better performance.

NOTE

For the most current information on performance tuning refer to the *QLogic OFED+ Host Software Release Notes*.

- **Use an IB MTU of 4096 bytes instead of 2048 bytes, if available, with the QLE7340, and QLE7342.** 4K MTU is enabled in the ib_qib driver by default. To change this setting for the driver, see “Changing the MTU Size” on page 3-20.
- **Make sure that write combining is enabled.** The x86 Page Attribute Table (PAT) mechanism that allocates Write Combining (WC) mappings for the PIO buffers has been added and is now the default. If PAT is unavailable or PAT initialization fails for some reason, the code will generate a message in the log and fall back to the MTRR mechanism. See Appendix F Write Combining for more information.
- **Check the PCIe bus width.** If slots have a smaller electrical width than mechanical width, lower than expected performance may occur. Use this command to check PCIe Bus width:

  ```
  $ ipath_control -iv
  ```

  This command also shows the link speed.
- **Experiment with non-default CPU affinity while running single-process-per-node latency or bandwidth benchmarks.** Latency may be slightly lower when using different CPUs (cores) from the default. On some chipsets, bandwidth may be higher when run from a non-default CPU or core. For the MPI being used, look at its documentation to see how to force a benchmark to run with a different CPU affinity than the default. With OFED micro benchmarks such as from the qperf or perftest suites, taskset will work for setting CPU affinity.

Remove Unneeded Services

The cluster administrator can enhance application performance by minimizing the set of system services running on the compute nodes. Since these are presumed to be specialized computing appliances, they do not need many of the service daemons normally running on a general Linux computer.

Following are several groups constituting a minimal necessary set of services. These are all services controlled by `chkconfig`. To see the list of services that are enabled, use the command:

```
$ /sbin/chkconfig --list | grep -w on
```

Basic network services are:

- `network`
- `ntpd`
- `syslog`
- `xinetd`

- `sshd`

For system housekeeping, use:

- `anacron`
- `atd`

- `crond`

If you are using Network File System (NFS) or yellow pages (yp) passwords:

- `rpcidmapd`
- `ypbind`
- `portmap`
- `nfs`
- `nfslock`

- `autofs`

To watch for disk problems, use:

- `smartd`

- `readahead`

The service comprising the ib_qib driver and SMA is:

- `openibd`

Other services may be required by your batch queuing system or user community.

If your system is running the daemon `irqbalance`, QLogic recommends turning it off. Disabling `irqbalance` will enable more consistent performance with programs that use interrupts. Use this command:

```
#  /sbin/chkconfig irqbalance off
```

See "Erratic Performance" on page D-10 for more information.

Host Environment Setup for MPI

After the QLogic OFED+ Host software and the GNU (GCC) compilers have been installed on all the nodes, the host environment can be set up for running MPI programs.

Configuring for ssh

Running MPI programs with the command `mpirun` on an IB cluster depends, by default, on secure shell `ssh` to launch node programs on the nodes.

To use `ssh`, you must have generated Rivest, Shamir, Adleman (RSA) or Digital Signal Algorithm (DSA) keys, public and private. The public keys must be distributed and stored on all the compute nodes so that connections to the remote machines can be established without supplying a password.

You or your administrator must set up the `ssh` keys and associated files on the cluster. There are two methods for setting up ssh on your cluster. The first method, the `shosts.equiv` mechanism, is typically set up by the cluster administrator. The second method, using `ssh-agent`, is more easily accomplished by an individual user.

NOTE

- `rsh` can be used instead of `ssh`. To use `rsh`, set the environment variable `MPI_SHELL=rsh`. See "Environment Variables" on page 4-18 for information on setting environment variables. Also see "Shell Options" on page A-6 for information on setting shell options in `mpirun`.
- `rsh` has a limit on the number of concurrent connections it can have, typically 255, which may limit its use on larger clusters.

Configuring `ssh` and `sshd` Using `shosts.equiv`

This section describes how the cluster administrator can set up `ssh` and `sshd` through the `shosts.equiv` mechanism. This method is recommended, provided that your cluster is behind a firewall and accessible only to trusted users.

“Configuring for ssh Using ssh-agent” on page 3-43 shows how an individual user can accomplish the same thing using `ssh-agent`.

The example in this section assumes the following:

- Both the cluster nodes and the front end system are running the `openssh` package as distributed in current Linux systems.
- All cluster end users have accounts with the same account name on the front end and on each node, by using Network Information Service (NIS) or another means of distributing the password file.
- The front end used in this example is called `ip-fe`.
- Root or superuser access is required on `ip-fe` and on each node to configure `ssh`.
- `ssh`, including the host’s key, has already been configured on the system `ip-fe`. See the `sshd` and `ssh-keygen` man pages for more information.

To use `shosts.equiv` to configure `ssg` and `sshd`:

1. On the system `ip-fe` (the front end node), change the `/etc/ssh/ssh_config` file to allow host-based authentication. Specifically, this file must contain the following four lines, all set to `yes`. If the lines are already there but commented out (with an initial #), remove the #.

   ```
   RhostsAuthentication yes
   RhostsRSAAuthentication yes
   HostbasedAuthentication yes
   EnableSSHKeysign yes
   ```

2. On each of the IB node systems, create or edit the file `/etc/ssh/shosts.equiv`, adding the name of the front end system. Add the line:

   ```
   ip-fe
   ```

   Change the file to mode 600 when you are finished editing.

3. On each of the IB node systems, create or edit the file `/etc/ssh/ssh_known_hosts`. You will need to copy the contents of the file `/etc/ssh/ssh_host_dsa_key.pub` from `ip-fe` to this file (as a single line), and then edit that line to insert `ip-fe ssh-dss` at the beginning of the line. This is very similar to the standard `known_hosts` file for `ssh`. An example line might look like this (displayed as multiple lines, but a single line in the file):

```
ip-fe ssh-dss
AAzAB3NzaC1kc3MAAACBAPoyES6+Akk+z3RfCkEHCkmYuYzqL2+1nwo4LeTVW
pCD1QsvrYRmpsfwpzYLXiSJdZSA8hfePWmMfrkvAAk4ueN8L3ZT4QfCTwqvHV
vSctpibf8n
aUmzloovBndOX9TIHyP/Ljfzzep4wL17+5hr1AHXldzrmgeEKp6ect1wxAAAA
FQDR56dAKFA4WgAiRmUJailtLFp8swAAAIBB1yrhF5P0jO+vpSnZrvrHa0Ok+
Y9apeJp3sessee30NlqKbJqWj5DOoRejr2VfTxZROf8LKuOY8tD6I59I0vlcQ
812E5iw1GCZfNefBmWbegWVKFwGlNbqBnZK7kDRLSOKQtuhYbGPcrVlSjuVps
fWEju64FTqKEetA8l8QEgAAAIBNtPDDwdmXRvDyc0gvAm6lPOIsRLmgmdgKXT
GOZUZ0zwxSL7GP1nEyFk9wAxCrXv3xPKxQaezQKs+KL95FouJvJ4qrSxxHdd1
NYNR0DavEBVQgCaspgWvWQ8cL
0aUQmTbggLrtD9zETVU5PCgRlQL6I3Y5sCCHuO7/UvTH9nneCg==
```

   Change the file to mode 600 when you are finished editing.

4. On each node, the system file `/etc/ssh/sshd_config` must be edited, so that the following four lines are uncommented (no # at the start of the line) and set to yes. (These lines are usually there, but are commented out and set to no by default.)

```
RhostsAuthentication yes
RhostsRSAAuthentication yes
HostbasedAuthentication yes
PAMAuthenticationViaKbdInt yes
```

5. After creating or editing the three files in Steps 2, 3, and 4, `sshd` must be restarted on each system. If you are already logged in via `ssh` (or any other user is logged in via `ssh`), their sessions or programs will be terminated, so restart only on idle nodes. Type the following (as root) to notify `sshd` to use the new configuration files:

```
# killall -HUP sshd
```

NOTE

This command terminates all `ssh` sessions into that system. Run from the console, or have a way to log into the console in case of any problem.

At this point, any end user should be able to login to the `ip-fe` front end system and use `ssh` to login to any IB node without being prompted for a password or pass phrase.

Configuring for `ssh` Using `ssh-agent`

The `ssh-agent`, a daemon that caches decrypted private keys, can be used to store the keys. Use `ssh-add` to add your private keys to `ssh-agent`'s cache. When `ssh` establishes a new connection, it communicates with `ssh-agent` to acquire these keys, rather than prompting you for a passphrase.

The process is described in the following steps:

1. Create a key pair. Use the default file name, and be sure to enter a passphrase.

   ```
   $ ssh-keygen -t rsa
   ```

2. Enter a passphrase for your key pair when prompted. Note that the key agent does not survive X11 logout or system reboot:

   ```
   $ ssh-add
   ```

3. The following command tells `ssh` that your key pair should let you in:

   ```
   $ cat ~/.ssh/id_rsa.pub >> ~/.ssh/authorized_keys
   ```

   Edit the `~/.ssh/config` file so that it reads like this:

   ```
   Host*
   ForwardAgent   yes
   ForwardX11   yes
   CheckHostIP   no
   StrictHostKeyChecking   no
   ```

   This file forwards the key agent requests back to your desktop. When you log into a front end node, you can use `ssh` to compute nodes without passwords.

4. Follow your administrator's cluster policy for setting up `ssh-agent` on the machine where you will be running `ssh` commands. Alternatively, you can start the `ssh-agent` by adding the following line to your `~/.bash_profile` (or equivalent in another shell):

   ```
   eval `ssh-agent`
   ```

   Use back quotes rather than single quotes. Programs started in your login shell can then locate the `ssh-agent` and query it for keys.

5. Finally, test by logging into the front end node, and from the front end node to a compute node, as follows:

```
$ ssh frontend_node_name
$ ssh compute_node_name
```

For more information, see the man pages for `ssh(1)`, `ssh-keygen(1)`, `ssh-add(1)`, and `ssh-agent(1)`.

Process Limitation with `ssh`

Process limitation with `ssh` is primarily an issue when using the `mpirun` option `-distributed=off`. The default setting is now `-distributed=on`; therefore, in most cases, `ssh` process limitations will not be encountered. This limitation for the `-distributed=off` case is described in the following paragraph.

MPI jobs that use more than 10 processes per node may encounter an `ssh` throttling mechanism that limits the amount of concurrent per-node connections to 10. If you need to use more processes, you or your system administrator must increase the value of `MaxStartups` in your `/etc/ssh/sshd_config` file.

Checking Cluster and Software Status

`ipath_control`

IB status, link speed, and PCIe bus width can be checked by running the program `ipath_control`. Sample usage and output are as follows:

```
$ ipath_control -iv
QLogic OFED.VERSION yyyy_mm_dd.hh_mm_ss
0: Version: ChipABI VERSION, InfiniPath_QLE7340, InfiniPath1
VERSION, SW Compat 2
0: Serial: RIB0935M31511 LocalBus: PCIe,5000MHz,x8
0,1: Status: 0xe1 Initted Present IB_link_up IB_configured
0,1: LID=0x23 GUID=0011:7500:005a:6ad0
0,1: HRTBT:Auto LINK:40 Gb/sec (4X QDR)
```

iba_opp_query

iba_opp_query is used to check the operation of the Distributed SA. You can run it from any node where the Distributed SA is installed and running, to verify that the replica on that node is working correctly. See “iba_opp_query” on page G-4 for detailed usage information.

```
# iba_opp_query --slid 0x31 --dlid 0x75 --sid 0x107
Query Parameters:
    resv1                   0x0000000000000107
    dgid                    ::
    sgid                    ::
    dlid                    0x75
    slid                    0x31
    hop                     0x0
    flow                    0x0
    tclass                  0x0
    num_path                0x0
    pkey                    0x0
    qos_class               0x0
    sl                      0x0
    mtu                     0x0
    rate                    0x0
    pkt_life                0x0
    preference              0x0
    resv2                   0x0
    resv3                   0x0
Using HCA qib0
Result:
    resv1                   0x0000000000000107
    dgid                    fe80::11:7500:79:e54a
    sgid                    fe80::11:7500:79:e416
    dlid                    0x75
    slid                    0x31
    hop                     0x0
    flow                    0x0
    tclass                  0x0
    num_path                0x0
    pkey                    0xffff
    qos_class               0x0
    sl                      0x1
```

```
mtu                     0x4
rate                    0x6
pkt_life                0x10
preference              0x0
resv2                   0x0
resv3                   0x0
```

ibstatus

Another useful program is `ibstatus` that reports on the status of the local HCAs. Sample usage and output are as follows:

```
$ ibstatus
Infiniband device 'qib0' port 1 status:
        default gid:
fe80:0000:0000:0000:0011:7500:005a:6ad0
        base lid:        0x23
        sm lid:          0x108
        state:           4: ACTIVE
        phys state:      5: LinkUp
        rate:            40 Gb/sec (4X QDR)
        link_layer:      IB
```

ibv_devinfo

ibv_devinfo queries RDMA devices. Use the -v option to see more information. Sample usage:

```
$ ibv_devinfo
hca_id: qib0
        fw_ver:                         0.0.0
        node_guid:                      0011:7500:00ff:89a6
        sys_image_guid:                 0011:7500:00ff:89a6
        vendor_id:                      0x1175
        vendor_part_id:                 29216
        hw_ver:                         0x2
        board_id:                       InfiniPath_QLE7280
        phys_port_cnt:                  1
                port:   1
                        state:                  PORT_ACTIVE (4)
                        max_mtu:                4096 (5)
                        active_mtu:             4096 (5)
                        sm_lid:                 1
                        port_lid:               31
                        port_lmc:               0x00
```

ipath_checkout

ipath_checkout is a bash script that verifies that the installation is correct and that all the nodes of the network are functioning and mutually connected by the IB fabric. It must be run on a front end node, and requires specification of a nodefile. For example:

```
$ ipath_checkout [options] nodefile
```

The nodefile lists the hostnames of the nodes of the cluster, one hostname per line. The format of nodefile is as follows:

```
hostname1
hostname2
...
```

For more information on these programs, see “ipath_control” on page G-27, “ibstatus” on page G-22, and “ipath_checkout” on page G-25.

4 Running MPI on QLogic Adapters

This section provides information on using the Message-Passing Interface (MPI) on QLogic IB HCAs. Examples are provided for setting up the user environment, and for compiling and running MPI programs.

Introduction

The MPI standard is a message-passing library or collection of routines used in distributed-memory parallel programming. It is used in data exchange and task synchronization between processes. The goal of MPI is to provide portability and efficient implementation across different platforms and architectures.

MPIs Packaged with QLogic OFED+

The high-performance open-source MPIs packaged with QLogic OFED+ include: Open MPI version 1.4.3, Ohio State University MVAPICH version 1.2, and MVAPICH2 version 1.7. These MPIs are offered in versions built with the high-performance Performance Scaled Messaging (PSM) interface and versions built run over IB Verbs. There are also the commercial MPIs which are not packaged with QOFED+, Intel MPI and Platform MPI, which both make use of the PSM application programming interface (API) and can both run over IB Verbs or over user direct access programming library (uDAPL), which uses IB Verbs. For more information on other MPIs, see Section 5 Using Other MPIs.

Open MPI

Open MPI is an open source MPI-2 implementation from the Open MPI Project. Pre-compiled versions of Open MPI version 1.4.3 that run over PSM and are built with the GCC, PGI, and Intel compilers are available with the QLogic download. Open MPI that runs over Verbs is also available.

Open MPI can be managed with the mpi-selector utility, as described in “Managing MVAPICH, and MVAPICH2 with the mpi-selector Utility” on page 5-5.

Installation

Follow the instructions in the *QLogic Fabric Software Installation Guide* for installing Open MPI.

Newer versions of Open MPI released after this QLogic OFED+ release will not be supported (refer to the *OFED+ Host Software Release Notes* for version numbers). QLogic does not recommend installing any newer versions of Open MPI. If a newer version is required it can be found on the Open MPI web site (http://www.open-mpi.org/) and installed after QLogic OFED+ has been installed.

Setup

When using the `mpi-selector` tool, the necessary `$PATH` and `$LD_LIBRARY_PATH` setup is done.

When not using the `mpi-selector` tool, put the Open MPI installation directory in the PATH by adding the following to `PATH`:

```
$mpi_home/bin
```

Where *`$mpi_home`* is the directory path where Open MPI is installed.

Compiling Open MPI Applications

QLogic recommends that you use the included wrapper scripts that invoke the underlying compiler (see Table 4-1).

Table 4-1. Open MPI Wrapper Scripts

Wrapper Script Name	Language
`mpicc`	C
`mpiCC, mpicxx, or mpic++`	C++
`mpif77`	Fortran 77
`mpif90`	Fortran 90

To compile your program in C, type the following:

```
$ mpicc mpi_app_name.c -o mpi_app_name
```

These scripts all provide the command line options listed in Table 4-2.

Table 4-2. Command Line Options for Scripts

Command	Meaning
`man mpicc (mpif90, mpicxx, etc.)`	Provides help
`-showme`	Lists each of the compiling and linking commands that would be called without actually invoking the underlying compiler
`-showme:compile`	Shows the compile-time flags that would be supplied to the compiler
`-showme:link`	Shows the linker flags that would be supplied to the compiler for the link phase.

These wrapper scripts pass most options on to the underlying compiler. Use the documentation for the underlying compiler (gcc, icc, pgcc, etc.) to determine what options to use for your application.

QLogic strongly encourages using the wrapper compilers instead of attempting to link to the Open MPI libraries manually. This allows the specific implementation of Open MPI to change without forcing changes to linker directives in users' Makefiles.

Create the `mpihosts` File

Create an MPI hosts file in the same working directory where Open MPI is installed. The MPI hosts file contains the host names of the nodes in your cluster that run the examples, with one host name per line. Name this file `mpihosts`. The contents can be in the following format:

More details on the `mpihosts` file can be found in “mpihosts File Details” on page 4-12.

Running Open MPI Applications

The Open MPI choices available from `mpi-selector --list` are:

- openmpi_gcc-1.4.3
- openmpi_gcc_qlc-1.4.3
- openmpi_intel_qlc-1.4.3
- openmpi_pgi_qlc-1.4.3.

The first choice will use verbs by default, and any with the `_qlc` string will use PSM by default. If you chose `openmpi_gcc_qlc-1.4.3`, for example, then the following simple mpirun command would run using PSM:

```
$ mpirun -np 4 -machinefile mpihosts mpi_app_name
```

To run over IB Verbs instead of the default PSM transport in `openmpi_gcc_qlc-1.4.3`, use this `mpirun` command line:

```
$ mpirun -np 4 -machinefile mpihosts --mca btl sm --mca btl
openib,self --mca mtl ^psm mpi_app_name
```

The following command enables shared memory:

```
--mca btl sm
```

The following command enables openib transport and communication to self:

```
--mca btl openib, self
```

The following command disables PSM transport:

```
--mca mtl ^psm
```

In these commands, `btl` stands for *byte transport layer* and `mtl` for *matching transport layer*.

PSM transport works in terms of MPI messages. OpenIB transport works in terms of byte streams.

Alternatively, you can use Open MPI with a sockets transport running over IPoIB, for example:

```
$ mpirun -np 4 -machinefile mpihosts --mca btl sm --mca btl
tcp,self --mca btl_tcp_if_exclude eth0 --mca
btl_tcp_if_include ib0 --mca mtl ^psm mpi_app_name
```

Note that `eth0` and `psm` are excluded, while `ib0` is included. These instructions may need to be adjusted for your interface names.

Note that in Open MPI, `machinefile` is also known as the `hostfile`.

Further Information on Open MPI

For more information about Open MPI, see:

http://www.open-mpi.org/

http://www.open-mpi.org/faq

Configuring MPI Programs for Open MPI

When configuring an MPI program (generating header files and/or Makefiles) for Open MPI, you usually need to specify `mpicc`, `mpicxx`, and so on as the compiler, rather than `gcc`, `g++`, etc.

Specifying the compiler is typically done with commands similar to the following, assuming that you are using `sh` or `bash` as the shell:

```
$ export CC=mpicc
$ export CXX=mpicxx
$ export F77=mpif77
$ export F90=mpif90
```

The shell variables will vary with the program being configured. The following examples show frequently used variable names. If you use `csh`, use commands similar to the following:

```
$ setenv CC mpicc
```

You may need to pass arguments to `configure` directly, for example:

```
$ ./configure -cc=mpicc -fc=mpif77 -c++=mpicxx
-c++linker=mpicxx
```

You may also need to edit a Makefile to achieve this result, adding lines similar to:

```
CC=mpicc
F77=mpif77
F90=mpif90
CXX=mpicxx
```

In some cases, the configuration process may specify the linker. QLogic recommends that the linker be specified as `mpicc`, `mpif90`, etc. in these cases. This specification automatically includes the correct flags and libraries, rather than trying to configure to pass the flags and libraries explicitly. For example:

```
LD=mpif90
```

These scripts pass appropriate options to the various compiler passes to include header files, required libraries, etc. While the same effect can be achieved by passing the arguments explicitly as flags, the required arguments may vary from release to release, so it is good practice to use the provided scripts.

To Use Another Compiler

Open MPI and all other MPIs that run on InfiniBand®, support a number of compilers, in addition to the default GNU Compiler Collection (GCC, including gcc, g++ and gfortran) versions 3.3 and later. These include the PGI 8.0, through 11.9; and Intel 9.*x*, 10.1, 11.*x*, and 12.*x*.

The easiest way to use other compilers with any MPI that comes with QLogic OFED+ is to use `mpi-selector` to change the selected MPI/compiler combination, see “Managing MVAPICH, and MVAPICH2 with the mpi-selector Utility” on page 5-5.

These compilers can be invoked on the command line by passing options to the wrapper scripts. Command line options override environment variables, if set.

Tables 4-3 and 4-4 show the options for each of the compilers.

In each case, `.....` stands for the remaining options to the `mpicxx` script, the options to the compiler in question, and the names of the files that it operates.

Table 4-3. Intel

Compiler	Command
C	`$ mpicc -cc=icc .....`
C++	`$ mpicc -CC=icpc`
Fortran 77	`$ mpif77 -fc=ifort .....`
Fortran 90/95	`$ mpif90 -f90=ifort .....` `$ mpif95 -f95=ifort .....`

Table 4-4. Portland Group (PGI)

Compiler	Command
C	`mpicc -cc=pgcc .....`
C++	`mpicc -CC=pgCC`
Fortran 77	`mpif77 -fc=pgf77 .....`
Fortran 90/95	`mpif90 -f90=pgf90 .....` `mpif95 -f95=pgf95 .....`

Also, use `mpif77`, `mpif90`, or `mpif95` for linking; otherwise, `.true.` may have the wrong value.

If you are not using the provided scripts for linking, link a sample program using the `-show` option as a test (without the actual build) to see what libraries to add to your link line. Some examples of the using the PGI compilers follow.

For Fortran 90 programs:

```
$ mpif90 -f90=pgf90 -show pi3f90.f90 -o pi3f90
pgf90 -I/usr/include/mpich/pgi5/x86_64 -c -I/usr/include
pi3f90.f90 -c
pgf90 pi3f90.o -o pi3f90 -lmpichf90 -lmpich
-lmpichabiglue_pgi5
```

Fortran 95 programs will be similar to the above.

For C programs:

```
$ mpicc -cc=pgcc -show cpi.c
pgcc -c cpi.c
pgcc cpi.o -lmpich -lpgftnrtl -lmpichabiglue_pgi5
```

Compiler and Linker Variables

When you use environment variables (e.g., `$MPICH_CC`) to select the compiler `mpicc` (and others) will use, the scripts will also set the matching linker variable (for example, `$MPICH_CLINKER`), if it is not already set. When both the environment variable and command line options are used (`-cc=gcc`), the command line variable is used.

When both the compiler and linker variables are set, and they do not match for the compiler you are using, the MPI program may fail to link; or, if it links, it may not execute correctly.

Process Allocation

Normally MPI jobs are run with each node program (process) being associated with a dedicated QLogic IB adapter *hardware context* that is mapped to a CPU.

If the number of node programs is greater than the available number of hardware contexts, *software context sharing* increases the number of node programs that can be run. Each adapter supports four software contexts per hardware context, so up to four node programs (from the same MPI job) can share that hardware context. There is a small additional overhead for each shared context.

Table 4-5 shows the maximum number of contexts available for each adapter.

Table 4-5. Available Hardware and Software Contexts

Adapter	Available Hardware Contexts (same as number of supported CPUs)	Available Contexts when Software Context Sharing is Enabled
QLE7342/ QLE7340	16	64

The default hardware context/CPU mappings can be changed on the QDR IB Adapters (QLE734*x*). See "IB Hardware Contexts on the QDR IB Adapters" on page 4-8 for more details.

Context sharing is enabled by default. How the system behaves when context sharing is enabled or disabled is described in "Enabling and Disabling Software Context Sharing" on page 4-9.

When running a job in a batch system environment where multiple jobs may be running simultaneously, it is useful to restrict the number of IB contexts that are made available on each node of an MPI. See "Restricting IB Hardware Contexts in a Batch Environment" on page 4-10.

Errors that may occur with context sharing are covered in "Context Sharing Error Messages" on page 4-11.

There are multiple ways of specifying how processes are allocated. You can use the `mpihosts` file, the `-np` and `-ppn` options with `mpirun`, and the `MPI_NPROCS` and `PSM_SHAREDCONTEXTS_MAX` environment variables. How these all are set are covered later in this document.

IB Hardware Contexts on the QDR IB Adapters

On the QLE7340 and QLE7342 QDR adapters, adapter receive resources are statically partitioned across the IB contexts according to the number of IB contexts enabled. The following defaults are automatically set according to the number of online CPUs in the node:

For four or less CPUs: 6 (4 + 2)

For five to eight CPUs: 10 (8 + 2)

For nine or more CPUs: 18 (16 + 2)

The one additional context on QDR adapters are to support the kernel on each port.

Performance can be improved in some cases by disabling IB hardware contexts when they are not required so that the resources can be partitioned more effectively.

To disable this behavior, explicitly configure for the number you want to use with the `cfgctxts` module parameter in the modprobe configuration file (see "Affected Files" on page 3-37 for exact file name and location).

The maximum that can be set is 18 on QDR IB Adapters.

The driver must be restarted if this default is changed. See "Managing the ib_qib Driver" on page 3-21.

NOTE

In rare cases, setting contexts automatically on QDR IB Adapters can lead to sub-optimal performance where one or more IB hardware contexts have been disabled and a job is run that requires software context sharing. Since the algorithm ensures that there is at least one IB context per online CPU, this case occurs only if the CPUs are over-subscribed with processes (which is not normally recommended). In this case, it is best to override the default to use as many IB contexts as are available, which minimizes the amount of software context sharing required.

Enabling and Disabling Software Context Sharing

By default, context sharing is enabled; it can also be specifically disabled.

Context Sharing Enabled: The MPI library provides PSM the local process layout so that IB contexts available on each node can be shared if necessary; for example, when running more node programs than contexts. All PSM jobs assume that they can make use of all available IB contexts to satisfy the job requirement and try to give a context to each process.

When context sharing is enabled on a system with multiple QLogic IB adapter boards (units) and the `IPATH_UNIT` environment variable is set, the number of IB contexts made available to MPI jobs is restricted to the number of contexts available on that unit. When multiple IB devices are present, it restricts the use to a specific IB Adapter unit. By default, all configured units are used in round robin order.

Context Sharing Disabled: Each node program tries to obtain exclusive access to an IB hardware context. If no hardware contexts are available, the job aborts.

To explicitly disable context sharing, set this environment variable in one of the two following ways:

```
PSM_SHAREDCONTEXTS=0
PSM_SHAREDCONTEXTS=NO
```

The default value of PSM_SHAREDCONTEXTS is 1 (enabled).

Restricting IB Hardware Contexts in a Batch Environment

If required for resource sharing between multiple jobs in batch systems, you can restrict the number of IB hardware contexts that are made available on each node of an MPI job by setting that number in the PSM_SHAREDCONTEXTS_MAX or PSM_RANKS_PER_CONTEXT environment variables.

For example, if you are running two different jobs on nodes using a QDR IB HCA, set PSM_SHAREDCONTEXTS_MAX to 8 instead of the default 16. Each job would then have at most 8 of the 16 available hardware contexts. Both of the jobs that want to share a node would have to set PSM_SHAREDCONTEXTS_MAX=8.

NOTE

MPIs use different methods for propagating environment variables to the nodes used for the job; See Section 7 for examples. Open MPI will automatically propagate PSM environment variables.

Setting PSM_SHAREDCONTEXTS_MAX=8 as a clusterwide default would unnecessarily penalize nodes that are dedicated to running single jobs. QLogic recommends that a per-node setting, or some level of coordination with the job scheduler with setting the environment variable should be used.

The number of contexts can be explicitly configured with the cfgctxts module parameter. This will override the default settings based on the number of CPUs present on each node. See “IB Hardware Contexts on the QDR IB Adapters” on page 4-8.

PSM_RANKS_PER_CONTEXT provides an alternate way of specifying how PSM should use contexts. The variable is the number of ranks that will share each hardware context. The supported values are 1, 2, 3 and 4, where 1 is no context sharing, 2 is 2-way context sharing, 3 is 3-way context sharing and 4 is the maximum 4-way context sharing. The same value of PSM_RANKS_PER_CONTEXT must be used for all ranks on a node, and typically, you would use the same value for all nodes in that job. Either PSM_RANKS_PER_CONTEXT or PSM_SHAREDCONTEXTS_MAX would be used in a particular job, but not both. If both are used and the settings are incompatible, then PSM will report an error and the job will fail to start up.

Context Sharing Error Messages

The error message when the context limit is exceeded is:

```
No free InfiniPath contexts available on /dev/ipath
```

This message appears when the application starts.

Error messages related to contexts may also be generated by `ipath_checkout` or `mpirun`. For example:

```
PSM found 0 available contexts on InfiniPath device
```

The most likely cause is that the cluster has processes using all the available PSM contexts. Clean up these processes before restarting the job.

Running in Shared Memory Mode

Open MPI supports running exclusively in shared memory mode; no QLogic adapter is required for this mode of operation. This mode is used for running applications on a single node rather than on a cluster of nodes.

To add pre-built applications (benchmarks), add `/usr/mpi/gcc/openmpi-1.4.3-qlc/tests/osu_benchmarks-3.1.1` to your PATH (or if you installed the MPI in another location: add *`$MPI_HOME`*`/tests/osu_benchmarks-3.1.1` to your PATH).

To enable shared memory mode, use a single node in the `mpihosts` file. For example, if the file were named `onehost` and it is in the working directory, the following would be entered:

```
$ cat /tmp/onehost
idev-64 slots=8
```

Enabling the shared memory mode as previously described uses a feature of Open-MPI host files to list the number of slots, which is the number of possible MPI processes (aka ranks) that you want to run on the node. Typically this is set equal to the number of processor cores on the node. A hostfile with 8 lines containing 'idev-64' would function identically. You can use this hostfile and run:

```
$ mpirun -np=2 -hostfile onehost osu_latency
```

to measure MPI latency between two cores on the same host using shared-memory, or

```
$ mpirun -np=2 -hostfile onehost osu_bw
```

to measure MPI unidirectional bandwidth using shared memory.

`mpihosts` File Details

As noted in “Create the mpihosts File” on page 4-3, a `hostfile` (also called *machines file, nodefile*, or *hostsfile*) has been created in your current working directory. This file names the nodes that the node programs may run.

The two supported formats for the `hostfile` are:

```
hostname1
hostname2
...
```

or

```
hostname1 slots=process_count
hostname2 slots=process_count
...
```

In the first format, if the `-np` count (number of processes to spawn in the `mpirun` command) is greater than the number of lines in the machine file, the hostnames will be repeated (in order) as many times as necessary for the requested number of node programs.

Also in the first format, if the `-np` count is less than the number of lines in the machine file, `mpirun` still processes the entire file and tries to pack processes to use as few hosts as possible in the hostfile. This is a different behavior than MVAPICH or the no-longer-supported QLogic MPI.

In the second format, *`process_count`* can be different for each host, and is normally the number of available processors on the node. When not specified, the default value is one. The value of *`process_count`* determines how many node programs will be started on that host before using the next entry in the `hostfile` file. When the full `hostfile` is processed, and there are additional processes requested, processing starts again at the start of the file.

It is generally recommended to use the second format and various command line options to schedule the placement of processes to nodes and cores. For example, the `mpirun` option `-npernode` can be used to specify (similar to the Intel MPI option `-ppn`) how many processes should be scheduled on each node on each pass through the hostfile. In the case of nodes with 8 cores each, if the hostfile line is specified as `hostname1 slots=8 max-slots=8`, then Open MPI will assign a maximum of 8 processes to the node and there can be no over-subscription of the 8 cores.

There are several alternative ways of specifying the `hostfile`:

- The command line option `-hostfile` can be used as shown in the following command line:

  ```
  $mpirun -np n -hostfile mpihosts [other options]
  program-name
  ```

  or `-machinefile` is a synonym for `-hostfile`. In this case, if the named file cannot be opened, the MPI job fails.

  An alternate mechanism to **`-hostfile`** for specifying hosts is the `-H`, `-hosts`, or `--host` followed by a host list. The host list can follow one of the following examples:

  ```
  host-01, or
  host-01,host-02,host-04,host-06,host-07,host-08
  ```

- In the absence of the `-hostfile` option, the `-H` option, `mpirun` uses the file `./mpihosts`, if it exists.

If you are working in the context of a batch queuing system, it may provide a job submission script that generates an appropriate `mpihosts` file. More details about how to schedule processes to nodes with Open MPI refer to the Open MPI website:

http://www.open-mpi.org/faq/?category=running#mpirun-scheduling

Using Open MPI's `mpirun`

The script `mpirun` is a front end program that starts a parallel MPI job on a set of nodes in an IB cluster. `mpirun` may be run on any x86_64 machine inside or outside the cluster, as long as it is on a supported Linux distribution, and has TCP connectivity to all IB cluster machines to be used in a job.

The script starts, monitors, and terminates the node programs. `mpirun` uses `ssh` (secure shell) to log in to individual cluster machines and prints any messages that the node program prints on `stdout` or `stderr`, on the terminal where `mpirun` is invoked.

`The general syntax is:`

```
$ mpirun [mpirun_options...] program-name [program options]
```

`program-name` is usually the pathname to the executable MPI program. When the MPI program resides in the current directory and the current directory is not in your search path, then *`program-name`* must begin with '`./`', for example:

```
./program-name
```

Unless you want to run only one instance of the program, use the `-np` option, for example:

```
$ mpirun -np n [other options] program-name
```

This option spawns *n* instances of *program-name*. These instances are called *node programs*.

Generally, mpirun tries to distribute the specified number of processes evenly among the nodes listed in the hostfile. However, if the number of processes exceeds the number of nodes listed in the hostfile, then some nodes will be assigned more than one instance of the program.

Another command line option, -npernode, instructs mpirun to assign a fixed number *p* of node programs (processes) to each node, as it distributes *n* instances among the nodes:

```
$ mpirun -np n -npernode p -hostfile mpihosts program-name
```

This option overrides the slots=*process_count* specifications, if any, in the lines of the mpihosts file. As a general rule, mpirun distributes the *n* node programs among the nodes without exceeding, on any node, the maximum number of instances specified by the slots=*process_count* option. The value of the slots=*process_count* option is specified by either the -npernode command line option or in the mpihosts file.

Typically, the number of node programs should not be larger than the number of processor cores, at least not for compute-bound programs.

This option specifies the number of processes to spawn. If this option is not set, then environment variable MPI_NPROCS is checked. If MPI_NPROCS is not set, the default is to determine the number of processes based on the number of hosts in the hostfile or the list of hosts -H or --host.

```
-npernode processes-per-node
```

This option creates up to the specified number of processes per node.

Each node program is started as a process on one node. While a node program may fork child processes, the children themselves must not call MPI functions.

There are many more mpirun options for scheduling where the processes get assigned to nodes. See man mpirun for details.

mpirun monitors the parallel MPI job, terminating when all the node programs in that job exit normally, or if any of them terminates abnormally.

Killing the mpirun program kills all the processes in the job. Use CTRL+C to kill mpirun.

Console I/O in Open MPI Programs

Open MPI directs UNIX standard input to /dev/null on all processes except the MPI_COMM_WORLD rank 0 process. The MPI_COMM_WORLD rank 0 process inherits standard input from mpirun.

NOTE

The node that invoked mpirun need not be the same as the node where the `MPI_COMM_WORLD` rank `0` process resides. Open MPI handles the redirection of `mpirun's` standard input to the rank `0` process.

Open MPI directs UNIX standard output and error from remote nodes to the node that invoked `mpirun` and prints it on the standard output/error of `mpirun`. Local processes inherit the standard output/error of `mpirun` and transfer to it directly.

It is possible to redirect standard I/O for Open MPI applications by using the typical shell redirection procedure on `mpirun`.

```
$ mpirun -np 2 my_app < my_input > my_output
```

Note that in this example only the `MPI_COMM_WORLD` rank `0` process will receive the stream from `my_input` on stdin. The stdin on all the other nodes will be tied to `/dev/null`. However, the stdout from all nodes will be collected into the `my_output` file.

Environment for Node Programs

The following information can be found in the Open MPI man page and is repeated here for easy of use.

Remote Execution

Open MPI requires that the PATH environment variable be set to find executables on remote nodes (this is typically only necessary in rsh- or ssh-based environments -- batch/scheduled environments typically copy the current environment to the execution of remote jobs, so if the current environment has PATH and/or LD_LIBRARY_PATH set properly, the remote nodes will also have it set properly). If Open MPI was compiled with shared library support, it may also be necessary to have the LD_LIBRARY_PATH environment variable set on remote nodes as well (especially to find the shared libraries required to run user MPI applications).

It is not always desirable or possible to edit shell startup files to set PATH and/or LD_LIBRARY_PATH. The `--prefix` option is provided for some simple configurations where this is not possible.

The `--prefix` option takes a single argument: the base directory on the remote node where Open MPI is installed. Open MPI will use this directory to set the remote PATH and LD_LIBRARY_PATH before executing any Open MPI or user applications. This allows running Open MPI jobs without having pre-configured the PATH and LD_LIBRARY_PATH on the remote nodes.

Open MPI adds the base-name of the current node's `bindir` (the directory where Open MPI's executables are installed) to the prefix and uses that to set the PATH on the remote node. Similarly, Open MPI adds the base-name of the current node's `libdir` (the directory where Open MPI's libraries are installed) to the prefix and uses that to set the LD_LIBRARY_PATH on the remote node. For example:

```
Local bindir:  /local/node/directory/bin
Local libdir:  /local/node/directory/lib64
```

If the following command line is used:

```
% mpirun --prefix /remote/node/directory
```

Open MPI will add `/remote/node/directory/bin` to the PATH and `/remote/node/directory/lib64` to the D_LIBRARY_PATH on the remote node before attempting to execute anything.

Note that `--prefix` can be set on a per-context basis, allowing for different values for different nodes.

The `--prefix` option is not sufficient if the installation paths on the remote node are different than the local node (for example, if `/lib` is used on the local node but `/lib64` is used on the remote node), or if the installation paths are something other than a subdirectory under a common prefix.

Note that executing `mpirun` using an absolute pathname is equivalent to specifying `--prefix` without the last subdirectory in the absolute pathname to `mpirun`. For example:

```
% /usr/local/bin/mpirun ...
is equivalent to
% mpirun --prefix /usr/local
```

Exported Environment Variables

All environment variables that are named in the form OMPI_* will automatically be exported to new processes on the local and remote nodes. The `-x` option to `mpirun` can be used to export specific environment variables to the new processes. While the syntax of the `-x` option allows the definition of new variables. Note that the parser for this option is currently not very sophisticated, it does not understand quoted values. Users are advised to set variables in the environment and use `-x` to export them, not to define them.

Setting MCA Parameters

The `-mca` switch allows the passing of parameters to various Modular Component Architecture (MCA) modules. MCA modules have direct impact on MPI programs because they allow tunable parameters to be set at run time (such as which BTL communication device driver to use, what parameters to pass to that BTL, and so on.).

The `-mca` switch takes two arguments: *`key`* and *`value`*. The *`key`* argument generally specifies which MCA module will receive the value. For example, the *`key`* `btl` is used to select which BTL to be used for transporting MPI messages. The *`value`* argument is the value that is passed. For example:

```
mpirun -mca btl tcp,self -np 1 foo
```

Tells Open MPI to use the `tcp` and `self` BTLs, and to run a single copy of `foo` an allocated node.

```
mpirun -mca btl self -np 1 foo
```

Tells Open MPI to use the `self` BTL, and to run a single copy of `foo` an allocated node.

The `-mca` switch can be used multiple times to specify different *`key`* and/or *`value`* arguments. If the same *`key`* is specified more than once, the `values` are concatenated with a comma (",") separating them.

Note that the `-mca` switch is simply a shortcut for setting environment variables. The same effect may be accomplished by setting corresponding environment variables before running `mpirun`. The form of the environment variables that Open MPI sets is:

```
OMPI_MCA_key=value
```

Thus, the `-mca` switch overrides any previously set environment variables. The `-mca` settings similarly override MCA parameters set in these two files, which are searched (in order):

1. `$HOME/.openmpi/mca-params.conf`: The user-supplied set of values takes the highest precedence.
2. `$prefix/etc/openmpi-mca-params.conf`: The system-supplied set of values has a lower precedence.

Environment Variables

Table 4-6 contains a summary of the environment variables that are relevant to any PSM including Open MPI. Table 4-7 is more relevant for the MPI programmer or script writer, because these variables are only active after the mpirun command has been issued and while the MPI processes are active. Open MPI provides the environmental variables shown in Table 4-7 that will be defined on every MPI process. Open MPI guarantees that these variables will remain stable throughout future releases.

Table 4-6. Environment Variables Relevant for any PSM

Name	Description
`OMP_NUM_THREADS`	This variable is used by a compilers' OpenMP run-time library.Use this variable to adjust the split between MPI processes and OpenMP threads. Usually, the number of MPI processes (per node) times the number of OpenMP threads will be set to match the number of CPUs per node. An example would be a node with eight CPUs, running two MPI processes and four OpenMP threads. In this case, OMP_NUM_THREADS is set to 4. OMP_NUM_THREADS is on a per-node basis, so needs to be propagated to each node used in the job, in a way that your MPI supports.
`PSM_SHAREDCONTEXTS`	This variable overrides automatic context sharing behavior. `YES` is equivalent to `1`. **Default:** `1`
`PSM_SHAREDCONTEXTS_MAX`	This variable restricts the number of IB contexts that are made available on each node of an MPI job. Up to 16 set automatically based on number of CPUs on node
`PSM_DEVICES`	Set this variable to enable running in shared memory mode on a single node.. **Default:** `self,ipath`

Table 4-6. Environment Variables Relevant for any PSM (Continued)

Name	Description
IPATH_NO_CPUAFFINITY	When set to 1, the PSM library will skip trying to set processor affinity. This is also skipped if the processor affinity mask is set to a list smaller than the number of processors prior to MPI_Init() being called. Otherwise the initialization code sets cpu affinity in a way that optimizes cpu and memory locality and load. **Default:** Unset
IPATH_PORT	Specifies the port to use for the job, 1 or 2. Specifying 0 will autoselect IPATH_PORT. **Default**: Unset
IPATH_UNIT	This variable is for context sharing. When multiple IB devices are present, and the IPATH_UNIT environment variable is set, the number of IB contexts made available to MPI jobs will be restricted to the number of contexts available on that unit. By default, IPATH_UNIT is unset and contexts from all configured units are made available to MPI jobs in round robin order. **Default:** Unset
IPATH_HCA_SELECTION_ALG	This variable provides user-level support to specify HCA/port selection algorithm through the environment variable. The default option is a round robin that allocates MPI processes to the HCAs in an alternating or round robin fashion. The older mechanism option is packed that fills all contexts on one HCA before allocating from the next HCA. For example: In the case of using two single-port HCAs, the default or IPATH_HCA_SELECTION_ALG= Round Robin setting, will allow 2 or more MPI processes per node to use both HCAs and to achieve performance improvements compared to what can be achieved with one HCA. **Default**: Round Robin
IPATH_SL	Service Level for QDR Adapters, these are used to work with the switch's Vfabric feature. **Default**: Unset

Table 4-6. Environment Variables Relevant for any PSM (Continued)

Name	Description
LD_LIBRARY_PATH	This variable specifies the path to the run-time library. **Default:** Unset

Table 4-7. Environment Variables Relevant for Open MPI

Name	Description
OMPI_COMM_WORLD_SIZE	This environment variable selects the number of processes in this process' MPI Comm_World
OMPI_COMM_WORLD_RANK	This variable is used to select the MPI rank of this process
OMPI_COMM_WORLD_LOCAL_RANK	This environment variable selects the relative rank of this process on this node within it job. For example, if four processes in a job share a node, they will each be given a local rank ranging from 0 to 3.
OMPI_UNIVERSE_SIZE	This environment variable selects the number of process slots allocated to this job. Note that this may be different than the number of processes in the job.

Job Blocking in Case of Temporary IB Link Failures

By default, as controlled by `mpirun's` quiescence parameter `-q`, an MPI job is killed for quiescence in the event of an IB link failure (or unplugged cable). This quiescence timeout occurs under one of the following conditions:

- A remote rank's process cannot reply to out-of-band process checks.
- MPI is inactive on the IB link for more than 15 minutes.

To keep remote process checks but disable triggering quiescence for temporary IB link failures, use the `-disable-mpi-progress-check` option with a nonzero `-q` option. To disable quiescence triggering altogether, use `-q 0`. No matter how these options are used, link failures (temporary or other) are always logged to `syslog`.

If the link is down when the job starts and you want the job to continue blocking until the link comes up, use the `-t -1` option.

Open MPI and Hybrid MPI/OpenMP Applications

Open MPI supports hybrid MPI/OpenMP applications, provided that MPI routines are called only by the master OpenMP thread. This application is called the *funneled thread model*. Instead of `MPI_Init/MPI_INIT` (for C/C++ and Fortran respectively), the program can call `MPI_Init_thread/MPI_INIT_THREAD` to determine the level of thread support, and the value `MPI_THREAD_FUNNELED` will be returned.

To use this feature, the application must be compiled with both OpenMP and MPI code enabled. To do this, use the `-openmp` or `-mp` flag (depending on your compiler) on the `mpicc` compile line.

As mentioned previously, MPI routines can be called only by the master OpenMP thread. The hybrid executable is executed as usual using `mpirun`, but typically only one MPI process is run per node and the OpenMP library will create additional threads to utilize all CPUs on that node. If there are sufficient CPUs on a node, you may want to run multiple MPI processes and multiple OpenMP threads per node.

The number of OpenMP threads is typically controlled by the `OMP_NUM_THREADS` environment variable in the `.bashrc` file. (`OMP_NUM_THREADS` is used by other compilers' OpenMP products, but is not an Open MPI environment variable.) Use this variable to adjust the split between MPI processes and OpenMP threads. Usually, the number of MPI processes (per node) times the number of OpenMP threads will be set to match the number of CPUs per node. An example case would be a node with four CPUs, running one MPI process and four OpenMP threads. In this case, `OMP_NUM_THREADS` is set to four. `OMP_NUM_THREADS` is on a per-node basis.

See "Environment for Node Programs" on page 4-15 for information on setting environment variables.

NOTE

With Open MPI, and other PSM-enabled MPIs, you will typically want to turn off PSM's CPU affinity controls so that the OpenMP threads spawned by an MPI process are not constrained to stay on the CPU core of that process, causing over-subscription of that CPU. Accomplish this using the `IPATH_NO_CPUAFFINITY=1` setting as follows:

`OMP_NUM_THREADS=8` (typically set in the `~/.bashrc` file)

```
mprun -np 2 -H host1,host2 -x IPATH_NO_CPUAFFINITY=1
./hybrid_app
```

In this case, typically there would be 8 or more CPU cores on the host1 and host2 nodes, and this job would run on a total of 16 threads, 8 on each node. You can use 'top' and then '1' to monitor that load is distributed to 8 different CPU cores in this case.

[Both the OMP_NUM_THREADS and IPATH_NO_CPUAFFINITY can be set in `.bashrc` or both on the command line after `-x` options.]

When there are more threads than CPUs, both MPI and OpenMP performance can be significantly degraded due to over-subscription of the CPUs

Debugging MPI Programs

Debugging parallel programs is substantially more difficult than debugging serial programs. Thoroughly debugging the serial parts of your code before parallelizing is good programming practice.

MPI Errors

Almost all MPI routines (except `MPI_Wtime` and `MPI_Wtick`) return an error code; either as the function return value in C functions or as the last argument in a Fortran subroutine call. Before the value is returned, the current MPI error handler is called. By default, this error handler aborts the MPI job. Therefore, you can get information about MPI exceptions in your code by providing your own handler for `MPI_ERRORS_RETURN`. See the man page for the `MPI_Errhandler_set` for details.

See the standard MPI documentation referenced in Appendix H for details on the MPI error codes.

Using Debuggers

- See http://www.open-mpi.org/faq/?category=debugging for details on debugging with Open MPI.

NOTE

The TotalView® debugger can be used with the Open MPI supplied in this release. Consult the TotalView documentation for more information:

http://www.open-mpi.org/faq/?category=running#run-with-tv

5 Using Other MPIs

This section provides information on using other MPI implementations. Detailed information on using Open MPI is provided in Section 4, and will be covered in this Section in the context of choosing among multiple MPIs or in tables which compare the multiple MPIs available.

Introduction

Support for multiple high-performance MPI implementations has been added. Most implementations run over both PSM and OpenFabrics Verbs (see Table 5-1). To choose which MPI to use, use the `mpi-selector-menu` command, as described in "Managing MVAPICH, and MVAPICH2 with the mpi-selector Utility" on page 5-5.

Table 5-1. Other Supported MPI Implementations

MPI Implementation	Runs Over	Compiled With	Comments
Open MPI 1.4.3	PSM Verbs	GCC, Intel, PGI	Provides some MPI-2 functionality (one-sided operations and dynamic processes). Available as part of the QLogic download. Can be managed by `mpi-selector`.
MVAPICH version 1.2	PSM Verbs	GCC, Intel, PGI	Provides MPI-1 functionality. Available as part of the QLogic download. Can be managed by `mpi-selector`.
MVAPICH2 version 1.7	PSM Verbs	GCC, Intel, PGI	Provides MPI-2 Functionality. Can be managed by `MPI-Selector`.
Platform MPI 8	PSM Verbs	GCC (default)	Provides some MPI-2 functionality (one-sided operations). Available for purchase from Platform Computing (an IBM Company).
Intel MPI version 4.0	TMI/PSM, uDAPL	GCC (default)	Provides MPI-1 and MPI-2 functionality. Available for purchase from Intel.

Table 5-1. Other Supported MPI Implementations (Continued)

MPI Implementation	Runs Over	Compiled With	Comments

Table Notes

MVAPICH and Open MPI have been have been compiled for PSM to support the following versions of the compilers:

- (GNU) gcc 4.1.0
- (PGI) pgcc 9.0
- (Intel) icc 11.1

These MPI implementations run on multiple interconnects, and have their own mechanisms for selecting the interconnect that runs on. Basic information about using these MPIs is provided in this section. However, for more detailed information, see the documentation provided with the version of MPI that you want to use.

Installed Layout

By default, the MVAPICH, MVAPICH2, and Open MPI are installed in the following directory tree:

`/usr/mpi/`*`$compiler`*`/`*`$mpi-mpi_version`*

The QLogic-supplied MPIs precompiled with the GCC, PGI, and the Intel compilers will also have `-qlc` appended after the MPI version number.

For example:

`/usr/mpi/gcc/openmpi-`*`VERSION`*`-qlc`

If a prefixed installation location is used, `/usr` is replaced by *`$prefix`*.

The following examples assume that the default path for each MPI implementation to `mpirun` is:

`/usr/mpi/`*`$compiler`*`/`*`$mpi`*`/bin/mpirun`

Again, `/usr` may be replaced by *`$prefix`*. This path is sometimes referred to as *`$mpi_home`*`/bin/mpirun` in the following sections.

See the documentation for Intel MPI, and Platform MPI for their default installation directories.

Open MPI

Open MPI is an open source MPI-2 implementation from the Open MPI Project. Pre-compiled versions of Open MPI version 1.4.3 that run over PSM and are built with the GCC, PGI, and Intel compilers are available with the QLogic download.

Details on Open MPI operation are provided in Section 4.

MVAPICH

Pre-compiled versions of MVAPICH 1.2 built with the GNU, PGI, and Intel compilers, and that run over PSM, are available with the QLogic download.

MVAPICH that runs over Verbs and is pre-compiled with the GNU compiler is also available.

MVAPICH can be managed with the `mpi-selector` utility, as described in "Managing MVAPICH, and MVAPICH2 with the mpi-selector Utility" on page 5-5.

Compiling MVAPICH Applications

As with Open MPI, QLogic recommends that you use the included wrapper scripts that invoke the underlying compiler (see Table 5-2).

Table 5-2. MVAPICH Wrapper Scripts

Wrapper Script Name	Language
`mpicc`	C
`mpiCC`, `mpicxx`	C++
`mpif77`	Fortran 77
`mpif90`	Fortran 90

To compile your program in C, type:

```
$ mpicc mpi_app_name.c -o mpi_app_name
```

To check the default configuration for the installation, check the following file:

```
/usr/mpi/$compiler/$mpi/etc/mvapich.conf
```

Running MVAPICH Applications

By default, the MVAPICH shipped with the QLogic OFED+ and IFS (IFS), runs over PSM once it is installed.

Here is an example of a simple `mpirun` command running with four processes:

```
$ mpirun -np 4 -hostfile mpihosts mpi_app_name
```

Password-less `ssh` is used unless the `-rsh` option is added to the command line above.

Further Information on MVAPICH

For more information about MVAPICH, see:

http://mvapich.cse.ohio-state.edu/

MVAPICH2

Pre-compiled versions of MVAPICH2 1.7 built with the GNU, PGI, and Intel compilers, and that run over PSM, are available with the QLogic download.

MVAPICH2 that runs over Verbs and is pre-compiled with the GNU compiler is also available.

MVAPICH2 can be managed with the `mpi-selector` utility, as described in "Managing MVAPICH, and MVAPICH2 with the mpi-selector Utility" on page 5-5.

Compiling MVAPICH2 Applications

As with Open MPI, QLogic recommends that you use the included wrapper scripts that invoke the underlying compiler (see Table 5-3).

Table 5-3. MVAPICH Wrapper Scripts

Wrapper Script Name	Language
`mpicc`	C
`mpiCC`, `mpicxx`	C++
`mpif77`	Fortran 77
`mpif90`	Fortran 90

To compile your program in C, type:

```
$ mpicc mpi_app_name.c -o mpi_app_name
```

To check the default configuration for the installation, check the following file:

```
/usr/mpi/$compiler/$mpi/etc/mvapich.conf
```

Running MVAPICH2 Applications

By default, the MVAPICH2 options in mpi-selector with 'qlc' as part of their name run over PSM once it is installed.

Here is an example of a simple mpirun command running with four processes:

```
$ mpirun_rsh -np 4 -hostfile mpihosts ./mpi_app_name
```

Further Information on MVAPICH2

For more information about MVAPICH2, see:

http://mvapich.cse.ohio-state.edu/support/mvapich2-1.7-quick-start.html

or for more detail:

http://mvapich.cse.ohio-state.edu/support/mvapich2-1.7rc2_user_guide.pdf

Managing MVAPICH, and MVAPICH2 with the mpi-selector Utility

When multiple MPI implementations have been installed on the cluster, you can use the mpi-selector to switch between them. The MPIs that can be managed with the mpi-selector are:

- MVAPICH
- MVAPICH2

The mpi-selector is an OFED utility that is installed as a part of QLogic OFED+ 1.5.4. Its basic functions include:

- Listing available MPI implementations
- Setting a default MPI to use (per user or site wide)
- Unsetting a default MPI to use (per user or site wide)
- Querying the current default MPI in use

Following is an example for listing and selecting an MPI:

```
$ mpi-selector --list
  mpi-1.2.3
  mpi-3.4.5
$ mpi-selector --set mpi-3.4.5
```

The new default takes effect in the next shell that is started. See the mpi-selector man page for more information.

The example shell scripts `mpivars.sh` and `mpivars.csh`, for registering with `mpi-selector`, are provided as part of the `mpi-devel` RPM in `$prefix/share/mpich/mpi-selector-{intel, gnu, pgi}` directories.

For all non-GNU compilers that are installed outside standard Linux search paths, set up the paths so that compiler binaries and runtime libraries can be resolved. For example, set `LD_LIBRARY_PATH`, both in your local environment and in an rc file (such as `.mpirunrc`, `.bashrc`, or `.cshrc`), are invoked on remote nodes. See "Environment for Node Programs" on page 4-15 and "Compiler and Linker Variables" on page 4-7 for information on setting up the environment for information on setting the run-time library path.

NOTE

The Intel-compiled versions require that the Intel compiler be installed and that paths to the Intel compiler runtime libraries be resolvable from the user's environment. The version used is Intel 10.1.012.

Platform MPI 8

Platform MPI 8 (formerly HP–MPI) is a high performance, production–quality implementation of the Message Passing Interface (MPI), with full MPI-2 funcionality. Platform MPI 8 is distributed by over 30 commercial software vendors, so you may need to use it if you use certain HPC applications, even if you don't purchase the MPI separately.

Installation

Follow the instructions for downloading and installing Platform MPI 8 from the Platform Computing web site.

Setup

Edit two lines in the `hpmpi.conf` file as follows:

Change,

```
MPI_ICMOD_PSM__PSM_MAIN = "^ib_ipath"
```

to,

```
MPI_ICMOD_PSM__PSM_MAIN = "^"
```

Change,

```
MPI_ICMOD_PSM__PSM_PATH = "^ib_ipath"
```

to,

```
MPI_ICMOD_PSM__PSM_PATH = "^"
```

Compiling Platform MPI 8 Applications

As with Open MPI, QLogic recommends that you use the included wrapper scripts that invoke the underlying compiler (see Table 5-4).

Table 5-4. Platform MPI 8 Wrapper Scripts

Wrapper Script Name	Language
mpicc	C
mpiCC	C
mpi77	Fortran 77
mpif90	Fortran 90

To compile your program in C using the default compiler, type:

```
$ mpicc mpi_app_name.c -o mpi_app_name
```

Running Platform MPI 8 Applications

Here is an example of a simple `mpirun` command running with four processes, over PSM:

```
$ mpirun -np 4 -hostfile mpihosts -PSM mpi_app_name
```

To run over IB Verbs, type:

```
$ mpirun -np 4 -hostfile mpihosts -IBV mpi_app_name
```

To run over TCP (which could be IPoIB if the hostfile is setup for IPoIB interfaces), type:

```
$ mpirun -np 4 -hostfile mpihosts -TCP mpi_app_name
```

More Information on Platform MPI 8

For more information on Platform MPI 8, see the Platform Computing web site

Intel MPI

Intel MPI version 4.0 is supported with this release.

Installation

Follow the instructions for download and installation of Intel MPI from the Intel web site.

Setup

Intel MPI can be run over Tag Matching Interface (TMI)

The setup for Intel MPI is described in the following steps:

1. Make sure that the TMI psm provider is installed on every node and all nodes have the same version installed. The TMI is supplied with the Intel MPI distribution. It can be installed either with the QLogic OFED+ Host Software installation or using the `rpm files`. For example:

   ```
   $ rpm -qa | grep tmi
   tmi-1.0-1
   ```

2. Verify that there is a `/etc/tmi.conf` file. It should be installed when installing the TMI. The file `tmi.conf` contains a list of TMI psm providers. In particular it must contain an entry for the PSM provider in a form similar to:

   ```
   psm X.X libtmip_psm.so " " # Comments OK
   ```

Intel MPI can also be run over uDAPL, which uses IB Verbs. uDAPL is the user mode version of the Direct Access Provider Library (DAPL), and is provided as a part of the OFED packages. You will also have to have IPoIB configured.

The setup for Intel MPI is described in the following steps:

1. Make sure that DAPL 1.2 or 2.0 is installed on every node and all nodes have the same version installed. In this release they are called `compat-dapl`. Both versions are supplied with the OpenFabrics RPMs and are included in the QLogic OFED+ Host Software package. They can be installed either with the QLogic OFED+ Host Software installation or using the `rpm files` after the QLogic OFED+ Host Software tar file has been unpacked. For example:

   Using DAPL 1.2.

   ```
   $ rpm -qa | grep compat-dapl
   compat-dapl-1.2.12-1.x86_64.rpm
   compat-dapl-debuginfo-1.2.12-1.x86_64.rpm
   compat-dapl-devel-1.2.12-1.x86_64.rpm
   compat-dapl-devel-static-1.2.12-1.x86_64.rpm
   compat-dapl-utils-1.2.12-1.x86_64.rpm
   ```

Using DAPL 2.0.

```
$ rpm -qa | grep dapl
dapl-devel-static-2.0.19-1
compat-dapl-1.2.14-1
dapl-2.0.19-1
dapl-debuginfo-2.0.19-1
compat-dapl-devel-static-1.2.14-1
dapl-utils-2.0.19-1
compat-dapl-devel-1.2.14-1
dapl-devel-2.0.19-1
```

2. Verify that there is a `/etc/dat.conf` file. It should be installed by the `dapl-` RPM. The file `dat.conf` contains a list of interface adapters supported by uDAPL service providers. In particular, it must contain mapping entries for OpenIB-cma for `dapl 1.2.x` and `ofa-v2-ib` for `dapl 2.0.x`, in a form similar to this (each on one line):

   ```
   OpenIB-cma u1.2 nonthreadsafe default libdaplcma.so.1
   dapl.1.2 "ib0 0" ""
   ```

   and

   ```
   ofa-v2-ib0 u2.0 nonthreadsafe default libdaplofa.so.2
   dapl.2.0 "ib0 0" ""
   ```

3. On every node, type the following command (as a root user):

   ```
   # modprobe rdma_ucm
   ```

   To ensure that the module is loaded when the driver is loaded, add `RDMA_UCM_LOAD=yes` to the `/etc/infiniband/openib.conf` file. (Note that `rdma_cm` is also used, but it is loaded automatically.)

4. Bring up an IPoIB interface on every node, for example, `ib0`. See “Configuring IPoIB” on page 3-5 for more details on configuring IPoIB.

Intel MPI has different bin directories for 32-bit (`bin`) and 64-bit (`bin64`); 64-bit is the most commonly used.

To launch MPI jobs, the Intel installation directory must be included in `PATH` and `LD_LIBRARY_PATH`.

When using `sh` for launching MPI jobs, run the following command:

```
$ source <$prefix>/bin64/mpivars.sh
```

When using `csh` for launching MPI jobs, run the following command:

```
$ source <$prefix>/bin64/mpivars.csh
```

Substitute `bin` if using 32-bit.

Compiling Intel MPI Applications

As with Open MPI, QLogic recommended that you use the included wrapper scripts that invoke the underlying compiler. The default underlying compiler is GCC, including gfortran. Note that there are more compiler drivers (wrapper scripts) with Intel MPI than are listed here (see Table 5-5); check the Intel documentation for more information.

Table 5-5. Intel MPI Wrapper Scripts

Wrapper Script Name	Language
`mpicc`	C
`mpiCC`	C++
`mpif77`	Fortran 77
`mpif90`	Fortran 90
`mpiicc`	C (uses Intel C compiler)
`mpiicpc`	C++ (uses Intel C++ compiler)
`mpiifort`	Fortran 77/90 (uses Intel Fortran compiler)

To compile your program in C using the default compiler, type:

```
$ mpicc mpi_app_name.c -o mpi_app_name
```

To use the Intel compiler wrappers (`mpiicc`, `mpiicpc`, `mpiifort`), the Intel compilers must be installed and resolvable from the user's environment.

Running Intel MPI Applications

Here is an example of a simple `mpirun` command running with four processes:

```
$ mpirun -np 4 -f mpihosts mpi_app_name
```

For more information, follow the Intel MPI instructions for usage of `mpirun`, `mpdboot`, and `mpiexec` (mpirun is a wrapper script that invoked both `mpdboot` and `mpiexec`). Remember to use `-r ssh` with `mpdboot` if you use `ssh`.

Pass the following option to `mpirun` to select TMI:

```
-genv I_MPI_FABRICS tmi
```

Pass the following option to `mpirun` to select uDAPL:

uDAPL 1.2:

```
-genv I_MPI_DEVICE rdma:OpenIB-cma
```

uDAPL 2.0:

```
-genv I_MPI_DEVICE rdma:ofa-v2-ib
```

To help with debugging, you can add this option to the Intel `mpirun` command:

TMI:

```
-genv TMI_DEBUG 1
```

uDAPL:

```
-genv I_MPI_DEBUG 2
```

Further Information on Intel MPI

For more information on using Intel MPI, see: http://www.intel.com/

Improving Performance of Other MPIs Over IB Verbs

Performance of MPI applications when using an MPI implementation over IB Verbs can be improved by tuning the IB MTU size.

NOTE

No manual tuning is necessary for PSM-based MPIs, since the PSM layer determines the largest possible IB MTU for each source/destination path.

The maximum supported MTU size of IB adapter cards is 4K.

Support for 4K IB MTU requires switch support for 4K MTU. The method to set the IB MTU size varies by MPI implementation:

- Open MPI defaults to the lower of either the IB MTU size or switch MTU size.
- MVAPICH defaults to an IB MTU size of 1024 bytes. This can be over-ridden by setting an environment variable:

  ```
  $ export VIADEV_DEFAULT_MTU=MTU2048
  ```

  Valid values are `MTU256`, `MTU512`, `MTU1024`, `MTU2048` and `MTU4096`. This environment variable must be set for all processes in the MPI job. To do so, use `~/.bashrc` or use of `/usr/bin/env`.
- MVAPICH2 defaults to an IB MTU size of 2048 bytes, which should be sufficient for most applications.
- Platform MPI over IB Verbs automatically determines the IB MTU size.
- Intel MPI over uDAPL (which uses IB Verbs) automatically determines the IB MTU size.

6 SHMEM Description and Configuration

Overview

QLogic SHMEM is a user-level communications library for one-sided operations. It implements the SHMEM Application Programming Interface (API) and runs on the QLogic IB stack. The SHMEM API provides global distributed shared memory across a network of hosts. Details of the API implementation are included in an appendix.

SHMEM is quite distinct from local shared memory (often abbreviated as "shm" or even “shmem”). Local shared memory is the sharing of memory by processes on the same host running the same OS system image. SHMEM provides access to global shared memory distributed across a cluster. The SHMEM API is completely different from and unrelated to the standard System V Shared Memory API provided by UNIX operating systems.

Interoperability

QLogic SHMEM depends on the Performance Scaled Messaging (PSM) protocol layer, implemented as a user-space library. QLogic SHMEM is only available to run with QLogic IB HCAs.

Installation

SHMEM is packaged with the QLogic IFS or QLogic OFED+ Host software.Every node in the cluster must have a QLogic IB adapter card and be running RedHat Enterprise Linux (RHEL) 6, 6.1 or 6.2 OS. One or more Message Passing Interface (MPI) implementations are required and Performance Scaled Messaging (PSM) support must be enabled within the MPI. The following MPI Implementations are supported:

- Open MPI version 1.4.3 configured to include PSM support. This is provided by QLogic IFS and can be found in the following directories:

```
/usr/mpi/gcc/openmpi-1.4.3-qlc
/usr/mpi/intel/openmpi-1.4.3-qlc
/usr/mpi/pgi/openmpi-1.4.3-qlc
```

The `-qlc` suffix denotes that this is the QLogic PSM version.

- MVAPICH version 1.2.0 compiled for PSM. This is provided by QLogic IFS and can be found in the following directories:

  ```
  /usr/mpi/gcc/mvapich-1.2.0-qlc
  /usr/mpi/intel/mvapich-1.2.0-qlc
  /usr/mpi/pgi/mvapich-1.2.0-qlc
  ```

  The `-qlc` suffix denotes that this is the QLogic PSM version.

- MVAPICH2 version 1.7 compiled for PSM. This is provided by QLogic IFS and can be found in the following directory:

  ```
  /usr/mpi/gcc/mvapich2-1.7-qlc
  /usr/mpi/intel/mvapich2-1.7-qlc
  /usr/mpi/pgi/mvapich2-1.7-qlc
  ```

  The `-qlc` suffix denotes that this is the QLogic PSM version.

It is recommended that you match the compiler used to build the MPI implementation with the compiler that you are using to build your SHMEM application. For example, if you are using the Intel compilers to build your SHMEM application and wish to run with Open MPI then use the Intel build of the Open MPI library:

```
/usr/mpi/intel/openmpi-1.4.3-qlc
```

The following C compilers are supported:

- gcc (as provided by distro) in 64-bit mode
- Intel 12.1 C compiler in 64-bit mode
- PGI 11.7 C compiler in 64-bit mode

For more information or to perform and installation with SHMEM enabled refer to Section 4 of the *QLogic Fabric Software Installation Guide*.

By default QLogic SHMEM is installed with a prefix of /usr/shmem/qlogic into the following directory structure:

```
/usr/shmem/qlogic
/usr/shmem/qlogic/bin
/usr/shmem/qlogic/bin/mvapich
/usr/shmem/qlogic/bin/mvapich2
/usr/shmem/qlogic/bin/openmpi
/usr/shmem/qlogic/lib64
/usr/shmem/qlogic/lib64/mvapich
/usr/shmem/qlogic/lib64/mvapich2
/usr/shmem/qlogic/lib64/openmpi
/usr/shmem/qlogic/include
```

QLogic recommends that `/usr/shmem/qlogic/bin` is added onto your `$PATH`.

If it is not on your `$PATH`, then you will need to give full pathnamescd to find the `shmemrun` and `shmemcc` wrapper scripts.

NOTE

There are subdirectories inside of bin for each MPI that are supported. These contain SHMEM benchmark programs that are linked directly against the MPI libraries as well as the SHMEM libraries.

SHMEM Programs

Basic SHMEM Program

Following is an example of a basic SHMEM program:

```
% cat shmem_world.c
#include <shmem.h>
#include <stdio.h>

int main ()
{
  shmem_init();
  printf("Hello from PE %d out of %d\n", my_pe(), num_pes());
  return 0;
  }
```

NOTE

These instructions assume a standard SHMEM installation and that `/usr/shmem/qlogic/bin` has been added to the $PATH.

The `%` character in the previous example is used to indicate the shell prompt and is followed by a command. The program can be compiled and linked using the `shmemcc` wrapper script:

```
% shmemcc shmem_world.c -o shmem_world
```

The program can be run using the `shmemrun` wrapper script:

```
% shmemrun -m hosts -np 2 ./shmem_world
  Hello from PE 1 out of 2
  Hello from PE 0 out of 2
```

This script assumes a `hosts` file is available, containing the host names on which the program is run. The `-np` option is used to specify the number of processing elements (PEs) to be run (for example, 2).

Compiling SHMEM Programs

The `shmemcc` script is a wrapper script for the compilation of the SHMEM C programs. The main purpose of the script is to call the C compiler with additional options to specify the SHMEM include directory, the SHMEM library directory, and to appropriately link in the SHMEM library. The `shmemcc` script automatically determines the correct directories by finding them relative to its own location. The standard directory layout of the QLogic SHMEM software is assumed.

The default C compiler is gcc, and can be overridden by specifying a compiler with the `$SHMEM_CC` environment variable.

If the option `-show` is added to the `shmemcc` command, it displays the command line that would be used to invoke the C compiler, but the C compiler will not be invoked. All other arguments to `shmemcc` are passed through to the C compiler without modification.

The C compiler can be used directly without using `shmemcc`. In that case the user must add the following to the command line:

For compilations add the following option:

`-I` *$SHMEM_DIR*`/include`

For linkages add the following options:

```
-Wl,--export-dynamic,--allow-shlib-undefined

-L $SHMEM_DIR/lib64/default

-lqlogic_shmem
```

Where *$SHMEM_DIR* in both of the options denotes the top-level directory of the SHMEM installation, typically the directory is `/usr/shmem/qlogic`.

The `-L` option uses the default version of the SHMEM libraries. The default is actually a symbolic link to libraries built for a specific MPI implementation. However, this choice does not constrain the SHMEM binary, and it can be run over any of the supported MPIs.

NOTE

If the SHMEM RPM is installed with `--prefix=usr` then the `-I` option is not necessary since the header files are in system default locations. All of the linkage options are still required.

The rationale for the `-Wl, --export-dynamic, --allow-shlib-undefined` options are to prevent other library and symbol dependencies in the SHMEM library from percolating up into the application binaries. These symbols include those from the underlying MPI implementation. There is no need to couple the application binary to a particular MPI, and these symbols will be correctly resolved at run-time. The advantage of this approach is that SHMEM application binaries will be portable across different implementations of the QLogic SHMEM library, including portability over different underlying MPIs.

Running SHMEM Programs

Using shmemrun

The `shmemrun` script is a wrapper script for running SHMEM programs using `mpirun`. The main purpose of the script is to call `mpirun` with additional options to specify the SHMEM library directory so that its dynamic libraries can be resolved. The script detects which `mpirun` is being used and remaps some common `mpirun` options to present a convenient and consistent interface to SHMEM users. Additionally, it enables PSM support in the underlying `mpirun` if required, and auto-propagates PSM, IPATH and SHMEM environment variables to the MPI processes. The `shmemrun` script automatically determines the correct directories by finding them relative to its own location. The `shmemrun` script can only automatically determine the correct directories if the standard directory layout of the QLogic SHMEM software has not been changed.

By default `mpirun` is picked up from the path and is assumed to be called `mpirun`. Alternatively, the pathname of `mpirun` can be specified with the `$SHMEM_MPIRUN` environment variable. There is also support for integration with slurm (see Slurm Integration). The following mpirun commands are supported:

- Open MPI: `mpirun`
- MVAPICH: `mpirun` and `mpirun_rsh`
- MVAPICH2: `mpirun` and `mpirun_rsh`

If the `shmemrun` script is run with `-show` option, it shows that the command line was used to invoke `mpirun`, but will not invoke it. Options that specify the number of processes and the hosts file are mapped by `shmemrun` to options that are accepted by the underlying `mpirun`. The contents of the host file can be parsed and regenerated if necessary and options to propagate environment variables are provided. The rationale for this script is to allow you to use the familiar options from the `mpirun` chosen and the options will automatically be remapped as required for the actual `mpirun`. This makes it possible to write scripts that call `shmemrun` without exposing these details of the underlying `mpirun` command.

If the `shmemrun` script finds the special -- option while processing the option list, that option is deleted and subsequent options and command line arguments are passed through without any modification. Using this option is useful to prevent `shmemrun` from modifying options of the program that are being run.

Running programs without using shmemrun

If you do not wish to use this wrapper script, then you must arrange for the SHMEM libraries to be found at run time using `$LD_LIBRARY_PATH` or an equivalent mechanism, and ensure that PSM support is enabled in your MPI implementation. The libraries can be found at:

`$SHMEM_DIR``/lib64/`*`$MPI`*

Where *`$SHMEM_DIR`* denotes the top-level directory of the SHMEM installation, typically `/usr/shmem/qlogic`, and *`$MPI`* is your choice of MPI (one of `mvapich`, `mvapich2`, or `openmpi`).

Additionally, the PSM receive thread and back-trace must be disabled using the following commands:

```
export PSM_RCVTHREAD=0
export IPATH_NO_BACKTRACE=1
```

QLogic SHMEM Relationship with MPI

QLogic SHMEM requires the QLogic PSM layer to provide the network transport function and this runs exclusively on QLogic IB HCAs. It also requires a compatible MPI implementation (also running over PSM) to provide program start up and other miscellaneous services. The one-sided operations in QLogic SHMEM are not layered on top of MPI, however, and go directly to PSM to give low-latency, high-performance access to the HCA architecture.

Typical SHMEM programs are written using calls to the SHMEM API and do not use MPI calls. In this case the program binary generated by `shmemcc` contains references to the SHMEM dynamic library and no references at all to MPI libraries. These binaries are portable across all MPI implementations supported by QLogic SHMEM. This is true of the `get/put` micro-benchmarks provided by QLogic SHMEM. The desired MPI can be selected at run time simply by placing the desired `mpirun` on `$PATH`, or by using the `$SHMEM_MPIRUN` environment variable.

Alternatively, it is possible to write hybrid SHMEM/MPI programs that use features from both the SHMEM and MPI libraries. These programs must call `shmem_init()` to initialize the SHMEM library state. They may also use `MPI_Init()` and `MPI_Finalize()` if needed. There will be a direct one-to-one correspondence between the SHMEM and `MPI_COMM_WORLD` rank assignments:

`shmem_my_pe()` will match `MPI_Comm_rank()` on `MPI_COMM_WORLD`

`shmem_n_pes()` will match `MPI_Comm_size()` on `MPI_COMM_WORLD`

Hybrid SHMEM/MPI programs must be linked against SHMEM libraries and the correct MPI libraries. It is recommended that the implementation of the MPI wrapper script(s) (`mpicc`) is used for compilation and that additional options are specified to find the SHMEM include and library files. One approach is to set up the `shmemcc` wrapper script to use `mpicc` as its compiler using the environment variable setting:

```
export SHMEM_CC=mpicc
```

This setting needs to be adjusted if `mpicc` is not already on the $PATH. The generated binary has references to both SHMEM and MPI libraries and is specific to that MPI implementation. QLogic recommends that `shmemrun` is used to run the program. The user must ensure that the correct `mpirun` is picked up from `$PATH` or using the `$SHMEM_MPIRUN` environment variable.

Slurm Integration

QLogic SHMEM relies on an MPI implementation to provide a run-time environment for jobs. This includes job start-up, stdin/stdout/stderr routing, and other low performance control mechanisms. QLogic SHMEM programs are typically started using shmemrun which is a wrapper script around mpirun. The shmemrun script takes care of setting up the environment appropriately, and also provides a common command-line interface regardless of which underlying mpirun is used.

Integration of QLogic SHMEM with slurm comes from the slurm integration provided by the MPI implementation. The slurm web pages describe 3 approaches. Please refer to points 1, 2 and 3 on the following web-page:

https://computing.llnl.gov/linux/slurm/mpi_guide.html

Below are various options for integration of the QLogic SHMEM and slurm.

Full Integration

This approach fully integrates QLogic SHMEM start-up into `slurm` and is available when running over MVAPICH2. The SHMEM program is executed using `srun` directly. For example:

```
srun -N 16 shmem-test-world
```

To run a program on 16 nodes. `slurm` starts the processes using `slurmd` and provides communication initialization. The implementation typically relies on `slurm` provided a process management interface (PMI) library and the MPI implementation using that so that each MPI process can hook into `slurm`.

The user is responsible for setting up the environment appropriately. This includes adding QLogic SHMEM's library directory to LD_LIBRARY_PATH. See Running SHMEM Programs for more information on the environment setup.

Two-step Integration

This approach is integrated, but is performed in 2 steps to allocate the nodes and run the job. This is available when running over Open MPI. The run command is now:

```
salloc -N 16 shmemrun shmem-test-world
```

The `salloc` allocates 16 nodes and runs one copy of `shmemrun` on the first allocated node which then creates the SHMEM processes. `shmemrun` invokes `mpirun`, and `mpirun` determines the correct set of hosts and required number of processes based on the `slurm` allocation that it is running inside of. Since `shmemrun` is used in this approach there is no need for the user to set up the environment.

No Integration

This approach allows a job to be launched inside a `slurm` allocation but with no integration. This approach can be used for any supported MPI implementation. However, it requires that a wrapper script is used to generate the hosts file. `slurm` is used to allocate nodes for the job, and the job runs within that allocation but not under the control of the `slurm` daemon. One way to use this approach is:

```
salloc -N 16 shmemrun_wrapper shmem-test-world
```

Where *`shmemrun_wrapper`* is a user-provided wrapper script that creates a hosts file based on the current `slurm` allocation and simply invokes `mpirun` with the hosts file and other appropriate options. Note that ssh/rsh will be used for starting processes not `slurm`.

Sizing Global Shared Memory

SHMEM provides `shmalloc`, `shrealloc` and `shfree` calls to allocate and release memory using a symmetric heap. These functions are called collectively across the processing elements (PEs) so that the memory is managed symmetrically across them. The extent of the symmetric heap determines the amount of global shared memory per PE that is available to the application.

This is an important resource and this section discusses the mechanisms available to size it. Applications can access this memory in various ways and this maps into quite different access mechanisms:

- Accessing global shared memory on my PE: This is achieved by direct loads and stores to the memory.
- Accessing global shared memory on a PE on the same host: This is achieved by mapping the global shared memory using the local shared memory mechanisms (for example, System V shared memory) operating system and then accessing the memory by direct loads and stores. This means that each PE on a host needs to map the global shared memory of each other PE on that host. These accesses do not use the adapter and interconnect.
- Accessing global shared memory on a PE on a different host: This is achieved by sending put, get, and atomic requests across the interconnect.

NOTE

There is a connection between the sizing of the global shared memory and local shared memory because of the mechanism used for accessing global shared memory in a PE that happens to be on the same host.

The QLogic SHMEM library pre-allocates room in the virtual address space according to `$SHMEM_SHMALLOC_MAX_SIZE` (default of 4GB). It then populates this with enough pages to cover `$SHMEM_SHMALLOC_INIT_SIZE` (default 16MB). The global shared memory segment can then grow dynamically from its initial size up to its maximum size. If an allocation attempts to exceed the maximum size allocations are no longer guaranteed to succeed, and will fail if there is no room in the virtual memory space of the process following the global shared memory segment. Upon failure the call to `shmalloc` or `shrealloc` returns NULL. The only down-side of using a large maximum size is occupancy of virtual address space (48 bits for 64-bit processes is very plentiful), and set-up of page table entries by the OS. A reasonable limit is 4GB per process. One side-effect of this approach is that SHMEM programs consume a large amount of virtual memory when viewed with the "top" program. This is due to the large maximum size setting. The `RES` field of `top` indicates the actual amount of memory that is resident in memory (for example, in actual use).

If a SHMEM application program runs out of global shared memory, increase the value of `$SHMEM_SHMALLOC_MAX_SIZE`. The value of `$SHMEM_SHMALLOC_INIT_SIZE` can also be changed to pre-allocate more memory up front rather than dynamically.

By default QLogic SHMEM will use the same base address for the symmetric heap across all PEs in the job. This address can be changed using the `$SHMEM_SHMALLOC_BASE_ADDR` environment variable. It will be rounded up to the nearest multiple of the page size. The virtual address range specified by this base address and the maximum size must not clash with any other memory mapping. If any SHMEM process in a job has a memory mapping clash, the QLogic SHMEM library will fail during `shmem_init()`. With 64-bit programs, a large virtual address space (for example, 48 bits in many modern processors) and a reasonably homogeneous cluster, it is expected that such failures will be rare. The default value of `$SHMEM_SHMALLOC_BASE_ADDR` has been chosen to work on the supported distributions and processors. In the rare event of a failure, the value of `$SHMEM_SHMALLOC_BASE_ADDR` can be changed using the environment variable.

Alternatively, if `$SHMEM_SHMALLOC_BASE_ADDR` is specified as `0`, then each SHMEM process will independently choose its own base virtual address for the global shared memory segment. In this case, the values for a symmetric allocation using `shmalloc()` are no longer guaranteed to be identical across the PEs. The QLogic SHMEM implementation takes care of this asymmetry by using offsets relative to the base of the symmetric heap in its protocols. However, applications that interpret symmetric heap pointer values or exchange symmetric heap pointer values between PEs will not behave as expected.

It is possible for SHMEM to fail at start-up or while allocating global shared memory due to limits placed by the operating system on the amount of *local* shared memory that SHMEM can use. Since SHMEM programs can use very large amounts of memory this can exceed typical OS configurations. As long as there is sufficient physical memory for the program, the following steps can be used to solve local shared memory allocation problems:

- Check for low `ulimits` on memory:

  `ulimit -l` : max locked memory (important for PSM not SHMEM)

  `ulimit -v` : max virtual memory

- Check the contents of these *sysctl* variables:

  `sysctl kernel.shmmax` ; maximum size of a single shm allocation in bytes

  `sysctl kernel.shmall` ; maximum size of all shm allocations in “pages”

  `sysctl kernel.shmnmi` ; maximum number of shm segments

- Check the size of `/dev/shm`:

  ```
  df /dev/shm
  ```

- Check for stale files in `/dev/shm`:

  ```
  ls /dev/shm
  ```

If any of these checks indicate a problem, ask the cluster administrator to increase the limit.

Progress Model

QLogic SHMEM supports active and passive progress models. Active progress means that the PE must actively call into SHMEM for progress to be made on SHMEM one-sided operations. Passive progress means that progress on SHMEM one-sided operations can occur without the application needing to call into SHMEM. Active progress is the default mode of operation for QLogic SHMEM. Passive progress can be selected using an environment variable where required.

Active Progress

In the active progress mode SHMEM progress is achieved when the application calls into the SHMEM library. This approach is well matched to applications that call into SHMEM frequently, for example, to have a fine grained mix of SHMEM operations and computation. This mix is typical of many SHMEM applications. Applications that spend large amount of contiguous time in computation without calling SHMEM routines will cause SHMEM progress to be delayed for that period of time. Additionally, applications must not poll on locations waiting for puts to arrive without calling SHMEM, since progress will not occur and the program will hang. Instead, SHMEM applications should use one of the wait synchronization primitives provided by SHMEM. In active progress mode QLogic SHMEM will achieve full performance.

Passive Progress

In the passive progress mode SHMEM progress will continue to occur when the application calls into SHMEM, but can additionally occur in the background when the application is not calling into SHMEM. This is achieved using an additional progress thread per PE. The progress thread is provided by PSM and is scheduled at a relatively low frequency, typically 10 to 100 times a second. This thread will cause independent SHMEM progress where required, both on the initiator side and the target side of SHMEM operations. In this mode applications can poll on locations waiting for puts to arrive without calling SHMEM. Progress will be achieved in this case by the progress thread, though it will incur the scheduling latency for the progress thread which may have a significant impact on overall performance if this idiom is used frequently. The scheduling frequency of the PSM progress thread can be tuned as described in the Environment Variables section.

Other performance effects of using passive progress include the following:

- The progress thread consumes some CPU cycles, though this is low because the progress thread runs infrequently.
- The SHMEM library uses additional locks in its implementation to protect its data structures against concurrent updates from the PE thread and the progress thread. There is a slight additional cost in the performance critical path because of this locking. This cost is minimal because contention on the lock is very low (the progress thread runs infrequently) and because each progress thread runs on the same CPU core as the corresponding PE thread (giving good cache locality for the lock).

- SHMEM's long message protocol is disabled. This is because the long message protocol implementation does not support passive progress. The effect of disabling this is to reduce long message bandwidth to that which can be achieved with the short message protocol. There is no effect on the bandwidth for message sizes below the long message break-point, which is set to 16KB by default.

Active versus Passive Progress

It is expected that most applications will be run with QLogic SHMEM's active progress mode since this gives full performance. The passive progress mode will typically be used in the following circumstances:

- For applications that use a polling idiom that is incompatible with the active progress mode, and where the application programmer is unable or unwilling to recode to use the appropriate SHMEM wait primitive.
- For compliance to a SHMEM standard that has a passive progress requirement.

Environment Variables

Table 6-1 list the environment variables that are currently provided by the SHMEM run time library.

NOTE

The set of supported environment variables and their defaults may vary.from release to release.

Table 6-1. SHMEM Run Time Library Environment Variables

Environment Variable	Default	Description
`$SHMEM_SHMALLOC_INIT_SIZE`	16MB	Initial size of the global shared memory segment.
`$SHMEM_SHMALLOC_MAX_SIZE`	4GB	Maximum size of the global shared memory segment.

Table 6-1. SHMEM Run Time Library Environment Variables (Continued)

Environment Variable	Default	Description
$SHMEM_SHMALLOC_CHECK	on	Shared memory consistency checks set for 0 to disable and 1 to enable. These are good checks for correctness but degrade the performance of shmalloc() and shfree(). These routines are usually not important for benchmark performance, so for now the checks are turned on to catch bugs early.
$SHMEM_IDENTIFY		If set, each SHMEM process will print out the SHMEM identity string and the path to the SHMEM library file.
$SHMEM_GET_REQ_LIMIT	64	Maximum number of outstanding short get requests for this end-point for the short get protocol (0 means unlimited). Each short get request can be up to 2KB.
$SHMEM_GET_LONG_REQ_LIMIT	16	Maximum number of outstanding get requests for this end-point for the long get protocol (0 means unlimited).
$SHMEM_PUT_FRAG_LIMIT	4096	Maximum number of outstanding put fragments for this end-point for the short put protocol (0 means unlimited). Each short put fragment can be up to 2KB.
$SHMEM_PUT_LONG_FRAG_LIMIT	128	Maximum number of outstanding put fragment requests for this end-point for the long get protocol (0 means unlimited).
$SHMEM_GET_LONG_SIZE	8KB for non-blocking `gets` 32KB for blocking `gets`	Gets of this size and larger use the SHMEM long get message protocol. Note that the parameter only allows the size to be changed in unison for both non-blocking and blocking gets.
$SHMEM_PUT_LONG_SIZE	8KB for non-blocking puts 16KB for blocking puts	Puts of this size and larger use the SHMEM long put message protocol. Note that the parameter only allows the size to be changed in unison for both non-blocking and blocking puts.

Table 6-1. SHMEM Run Time Library Environment Variables (Continued)

Environment Variable	Default	Description
$SHMEM_PUT_REPLY_COMBINING_COUNT	8	Number of consecutive put replies on a flow to combine together into a single reply.

The command `shmemrun` automatically propagates `SHMEM*` environment variables from its own environment to all the SHMEM processes. This means that the environment variables can be simply setup in the front-end shell used to invoke `shmemrun`. The command `shmemrun` also has its own environment variables that are listed in Table 6-2.

Table 6-2. `shmemrun` Environment Variables

Environment Variable	Default	Description
$SHMEM_MPIRUN	`mpirun` from the PATH	Specifies where to find `mpirun`.
$SHMEMRUN_VERBOSE		Enables verbose output for `shmemrun`.
$SHMEMRUN_SLEEP		Specifies a sleep time (in seconds) after the job completes. This variable is intended for testing use.
$SHMEMRUN_TIMEOUT		Specifies a time-out value (in seconds). When the timeout value is reached, the `mpirun` is killed. This variable is intended for testing use.

Implementation Behavior

Some SHMEM properties are not fully specified by the SHMEM API specification. This section discusses the behavior for the QLogic SHMEM implementation.

For a put operation, these descriptions use the terms "local completion" and “remote completion”. Once a put is locally complete, the source buffer on the initiating PE is available for reuse. Until a put is locally complete the source buffer must not be modified since that buffer is in use for the put operation. A blocking put is locally complete immediately upon return from the put. A non-blocking put is not locally complete upon return from the put—different mechanisms are used for detecting local completion using either an explicit handle (use `shmem_test_nb()` or `shmem_wait_nb())` or a NULL handle (use `shmem_quiet()`). Once a put is remotely complete the destination buffer on the target PE is fully written and available for use. The mechanism provided by SHMEM for detecting remote completion are described below.

- `shmem_fence()` - This function ensures that all puts issued by this PE prior to the fence will become remotely visible before any puts issued by this PE after the fence. The call does not necessarily imply that any of the prior puts are actually remotely visible at the point of the fence, only that this ordering is guaranteed.

- `shmem_quiet()` - This function waits for remote completion of all puts issued by this PE prior to the quiet operation. Therefore, once the quiet operation returns, it is guaranteed that all those puts will be remotely visible to other PEs. This guarantee of remote completion applies to all puts—blocking puts, non-blocking puts with handles, and non-blocking puts with NULL handles. Additionally, this function additionally waits for local completion of non-blocking puts and non-blocking gets that were issued with a NULL handle.

- `shmem_test_nb()` and `shmem_wait_nb()` can be used to test and wait for local completion of a `non_blocking` operation. For a non-blocking put, this does not indicate whether remote completion has occurred.

Additional properties of the QLogic SHMEM implementation are:

- The QLogic SHMEM implementation makes no guarantees as to the ordering in which the bytes of a put operation are delivered into the remote memory. It is *not* a safe assumption to poll or read certain bytes of the put destination buffer (for example, the last 8 bytes) to look for a change in value and then infer that the entirety of the put has arrived. The correct mechanism for this is to use the `shmem_quiet()` operation to force remote completion, or to use the following type of sequence:

  - Initiator side:

    - Issue a batch of puts all unordered with respect to each other

    - `shmem_fence()`

- 8 byte put to a sync location

- Target side:
  - Wait for the sync location to be written
  - Now it is safe to make observations on all puts prior to fence

- `shmem_int_wait()`, `shmem_long_wait()`, `shmem_longlong_wait()`, `shmem_short_wait()`, `shmem_wait()`, `shmem_int_wait_until()`, `shmem_long_wait_until()`, `shmem_longlong_wait_until()`, `shmem_short_wait_until()`, `shmem_wait_until()` - These SHMEM operations are provided for waiting for a variable in local symmetric memory to change value due to an incoming put. In the active progress mode SHMEM applications must use these routines for this purpose, and not implement their own polling loop without SHMEM library calls. In the passive progress mode SHMEM application may use a polling loop without a SHMEM library call. However, performance will typically be substantially improved by using the SHMEM wait operation instead.

- `shmem_stack()` is implemented as a no-op since this is a distributed memory cluster architecture.

- `shmem_ptr`(void *target, int pe) returns the provided address, if the PE is my PE, otherwise NULL. This implementation is sufficient to conform to the SHMEM API and is appropriate for a distributed memory cluster architecture.

- `shmem_clear_cache_inv()`, `shmem_clear_cache_line_inv()`, `shmem_set_cache_inv()`, `shmem_set_cache_line_inv()`, `shmem_udcflush()`, and `shmem_udcflush_line()` are each implemented as a no-op since there is no global memory caching in this implementation.

- This SHMEM implementation allows remote access to variables that are in the symmetric heap and static data/read-only data sections only. It does not support static data sections in dynamically loaded libraries.

Application Programming Interface

Table 6-3 lists the provided SHMEM Application Programming Interface (API) calls and details any restrictions.

Table 6-3. SHMEM Application Programming Interface Calls

Operation	Calls
General Operations	shmem_init
	start_pes
	my_pe
	_my_pe
	shmem_my_pe
	num_pes
	_num_pes
	shmem_n_pes
Symmetric heap	shmalloc
	shmemalign
	shfree
	shrealloc
Contiguous Put Operations	shmem_short_p
	shmem_int_p
	shmem_long_p
	shmem_float_p
	shmem_double_p
	shmem_longlong_p
	shmem_longdouble_p
	shmem_char_put
	shmem_short_put
	shmem_double_put
	shmem_float_put
	shmem_int_put
	shmem_long_put
	shmem_longdouble_put
	shmem_longlong_put

Table 6-3. SHMEM Application Programming Interface Calls

Operation	Calls
	shmem_put
	shmem_put32
	shmem_put64
	shmem_put128
	shmem_putmem
Non-blocking Put Operations	shmem_double_put_nb
	shmem_float_put_nb
	shmem_int_put_nb
	shmem_long_put_nb
	shmem_longdouble_put_nb
	shmem_longlong_put_nb
	shmem_put_nb
	shmem_put32_nb
	shmem_put64_nb
	shmem_put128_nb
	shmem_putmem_nb
	shmem_short_put_nb
Strided Put Operations	shmem_double_iput
	shmem_float_iput
	shmem_int_iput
	shmem_iput
	shmem_iput32
	shmem_iput64
	shmem_iput128
	shmem_long_iput
	shmem_longdouble_iput
	shmem_longlong_iput

Table 6-3. SHMEM Application Programming Interface Calls

Operation	Calls
	shmem_short_iput
Indexed Put Operations	shmem_ixput
	shmem_ixput32
	shmem_ixput64
Put and Non-blocking Ordering, Flushing and Completion	shmem_fence
	shmem_quiet
	shmem_wait_nb
	shmem_test_nb
	shmem_poll_nb (same as shmem_test_nb, provided for compatibility)
Contiguous Get Operations	shmem_short_g
	shmem_int_g
	shmem_long_g
	shmem_float_g
	shmem_double_g
	shmem_longlong_g
	shmem_longdouble_g
	shmem_char_get
	shmem_short_get
	shmem_double_get
	shmem_float_get
	shmem_int_get
	shmem_long_get
	shmem_longdouble_get
	shmem_longlong_get
	shmem_get

Table 6-3. SHMEM Application Programming Interface Calls

Operation	Calls
	shmem_get32
	shmem_get64
	shmem_get128
	shmem_getmem
Non-blocking Get Operations	shmem_double_get_nb
	shmem_float_get_nb
	shmem_int_get_nb
	shmem_long_get_nb
	shmem_longdouble_get_nb
	shmem_longlong_get_nb
	shmem_short_get_nb
	shmem_get_nb
	shmem_get32_nb
	shmem_get64_nb
	shmem_get128_nb
	shmem_getmem_nb
Strided Get Operations	shmem_double_iget
	shmem_float_iget
	shmem_int_iget
	shmem_iget
	shmem_iget32
	shmem_iget64
	shmem_iget128
	shmem_long_iget
	shmem_longdouble_iget
	shmem_longlong_iget
	shmem_short_iget

Table 6-3. SHMEM Application Programming Interface Calls

Operation	Calls
Indexed Get Operations	shmem_ixget
	shmem_ixget32
	shmem_ixget64
Barriers	barrier
	shmem_barrier_all
	shmem_barrier
Broadcasts	shmem_broadcast
	shmem_broadcast32
	shmem_broadcast64
Concatenation	shmem_collect
	shmem_collect32
	shmem_collect64
	shmem_fcollect
	shmem_fcollect32
	shmem_fcollect64
Synchronization operations	shmem_int_wait
	shmem_long_wait
	shmem_longlong_wait
	shmem_short_wait
	shmem_wait
	shmem_int_wait_until
	shmem_long_wait_until
	shmem_longlong_wait_until
	shmem_short_wait_until
	shmem_wait_until
Atomic operations	shmem_double_swap
	shmem_float_swap

Table 6-3. SHMEM Application Programming Interface Calls

Operation	Calls
	shmem_short_swap
	shmem_int_swap
	shmem_long_swap
	shmem_longlong_swap
	shmem_swap
	shmem_short_cswap
	shmem_int_cswap
	shmem_long_cswap
	shmem_longlong_cswap
	shmem_short_mswap
	shmem_int_mswap
	shmem_long_mswap
	shmem_longlong_mswap
	shmem_short_inc
	shmem_int_inc
	shmem_long_inc
	shmem_longlong_inc
	shmem_short_add
	shmem_int_add
	shmem_long_add
	shmem_longlong_add
	shmem_short_finc
	shmem_int_finc
	shmem_long_finc
	shmem_longlong_finc
	shmem_short_fadd
	shmem_int_fadd

Table 6-3. SHMEM Application Programming Interface Calls

Operation	Calls
	shmem_long_fadd
	shmem_longlong_fadd
Reductions	shmem_int_and_to_all
	shmem_long_and_to_all
	shmem_longlong_and_to_all
	shmem_short_and_to_all
	shmem_int_or_to_all
	shmem_long_or_to_all
	shmem_longlong_or_to_all
	shmem_short_or_to_all
	shmem_int_xor_to_all
	shmem_long_xor_to_all
	shmem_longlong_xor_to_all
	shmem_short_xor_to_all
	shmem_double_min_to_all
	shmem_float_min_to_all
	shmem_int_min_to_all
	shmem_long_min_to_all
	shmem_longdouble_min_to_all
	shmem_longlong_min_to_all
	shmem_short_min_to_all
	shmem_double_max_to_all
	shmem_float_max_to_all
	shmem_int_max_to_all
	shmem_long_max_to_all
	shmem_longdouble_max_to_all
	shmem_longlong_max_to_all

Table 6-3. SHMEM Application Programming Interface Calls

Operation	Calls
	shmem_short_max_to_all
	shmem_complexd_sum_to_all complex collectives are not implemented
	shmem_complexf_sum_to_all complex collectives are not implemented
	shmem_double_sum_to_all
	shmem_float_sum_to_all
	shmem_int_sum_to_all
	shmem_long_sum_to_all
	shmem_longdouble_sum_to_all
	shmem_longlong_sum_to_all
	shmem_short_sum_to_all
	shmem_complexd_prod_to_all complex collectives are not implemented
	shmem_complexf_prod_to_all complex collectives are not implemented
	shmem_double_prod_to_all
	shmem_float_prod_to_all
	shmem_int_prod_to_all
	shmem_long_prod_to_all
	shmem_longdouble_prod_to_all
	shmem_longlong_prod_to_all
	shmem_short_prod_to_all
All-to-all (an extension beyond classic SHMEM)	shmem_alltoall
	shmem_alltoall32
	shmem_alltoall64
Locks	shmem_set_lock

Table 6-3. SHMEM Application Programming Interface Calls

Operation	Calls
	`shmem_clear_lock`
	`shmem_test_lock`
Events	`clear_event`
	`set_event`
	`wait_event`
	`test_event`
General Operations (for compatibility)	`globalexit` allows any process to abort the job
	`shmem_finalize` call to terminate the SHMEM library
	`shmem_pe_accessible` tests PE for accessibility
	`shmem_addr_accessible` test address on PE for accessibility
Cache Operations (for compatibility)	`shmem_clear_cache_inv` implemented as a no-op
	`shmem_clear_cache_line_inv` implemented as a no-op
	`shmem_set_cache_inv` implemented as a no-op
	`shmem_set_cache_line_inv` implemented as a no-op
	`shmem_udcflush` implemented as a no-op
	`shmem_udcflush_line` implemented as a no-op
Stack/Pointer Operations (for compatibility)	`shmem_stack` implemented as a no-op

Table 6-3. SHMEM Application Programming Interface Calls

Operation	Calls
	`shmem_ptr` returns the address if the PE is my PE, otherwise NULL

SHMEM Benchmark Programs

The following SHMEM micro-benchmark programs are included:

- `shmem-get-latency`: measures get latency
- `shmem-get-bw`: measures streaming get bandwidth (uni-directional)
- `shmem-get-bibw`: measures streaming get bandwidth (bi-directional)
- `shmem-put-latency`: measures put latency
- `shmem-put-bw`: measures streaming put bandwidth (uni-directional)
- `shmem-put-bibw`: measures streaming put bandwidth (bi-directional)

The programs can be used to measure round-trip get latency, one way put latency, get and put bandwidth, as well as get and put message rates.

The benchmarks must be run with an even number of processes. They are typically run on exactly two hosts with the processes equally-divided between them. The processes are split up into pairs, with one from each pair on either host and each pair is loaded with the desired traffic pattern. The benchmark automatically determines the correct mapping, regardless of the actual rank order of the processes and their mapping to the two hosts.

Alternatively, if the `-f` option is specified the benchmark is forced to use the rank order when arranging the communication pattern. In this mode and with `np` ranks, each rank `i` in (`0`, `np/2`) is paired with rank `(np / 2) + i`. For example, this mode can be used to test SHMEM performance within a single node.

The micro-benchmarks have the command line options shown in Table 6-4

Table 6-4. QLogic SHMEM micro-benchmarks options

Option	Description
`-a INT`	a `log2` of desired alignment for buffers (default = 12)

Table 6-4. QLogic SHMEM micro-benchmarks options

Option	Description
`-b INT`	batch size, number of concurrent operations (default = 64)
`-f`	force order for bifurcation of PEs based on rank order
`-h`	displays the help page
`-l INT`	set minimum message size (default = 2)
`-m INT`	sets the maximum message size (default = 4194304)

Additional SHMEM micro-benchmark programs are included to measure get and put performance with randomized PE selection and randomized target memory locations, all-to-all communication patterns using put, barrier and reduce:

QLogic SHMEM random access benchmark

`shmem-rand`: randomized `put`/`get` benchmark

This is actually a hybrid SHMEM/MPI code, so a binary is provided per supported MPI implementation. It has the following command line options:

Usage: `shmem-rand [options] [list of message sizes].`

Message sizes are specified in bytes (default = 8)

Options: See Table 6-5

Table 6-5. QLogic SHMEM random access benchmark options

Option	Description
`-a`	use automatic (NULL) handles for NB ops (default explicit handles)
`-b`	use a barrier every window
`-c INTEGER`	specify loop count (see also `-t`)
`-f`	fixed window size (default is scaled)
`-h`	displays the help page
`-l`	enable communication to local ranks
`-m INTEGER[K]`	memory size in MB (default = 8MB): or in KB with a K suffix
`-n`	use non-pipelined mode for NB ops (default pipelined)

Table 6-5. QLogic SHMEM random access benchmark options

Option	Description
`-o OP`	choose `OP` from `get`, `getnb`, `put`, `putnb`
-p	for blocking puts, no quiet every window (this is the default)
-q	for blocking puts, use quiet every window
`-r`	use ring pattern (default is random)
`-s`	enable communication to self
`-t FLOAT`	if the loop count is not given, run the test for this many seconds (default is 10s)
`-u`	run in uni-directional mode
`-v`	verbose mode (repeat for more verbose)
`-w INTEGER`	set the window size (default = 32)
`-x INTEGER`	window size limit (default = 16384)

QLogic SHMEM all-to-all benchmark

`shmem-alltoall`: all-to-all put benchmark

This is a hybrid SHMEM/MPI code, so a binary is provided per supported MPI implementation. It has the following command line options:

Usage: `/test/shmem-alltoall [options] [list of message sizes]`

Message sizes are specified in bytes (default 8)

Options: See Table 6-6

Table 6-6. QLogic SHMEM all-to-all benchmark options

Option	Description
`-a`	use automatic (NULL) handles for NB ops (default explicit handles)
`-c INTEGER`	specify loop count (see also `-t`)
`-f`	fixed window size (default is scaled)
`-h`	displays the help page

Table 6-6. QLogic SHMEM all-to-all benchmark options

Option	Description
`-l`	enable communication to local ranks (including self)
`-m INTEGER[K]`	memory size in MB (default = 8MB): or in KB with a K suffix
`-n`	use non-pipelined mode for NB ops (default pipelined)
`-o OP`	choose `OP` from put, or `putnb`
-p INTEGER	offset for all-to-all schedule (default 1, usually set to ppn)
`-r`	randomize all-to-all schedule
`-s`	enable communication to self
`-t FLOAT`	if the loop count is not given, run the test for this many seconds (default is 10s)
`-v`	verbose mode (repeat for more verbose)
`-w INTEGER`	set the window size (default = 32)
`-x INTEGER`	window size limit (default = 16384)

QLogic SHMEM barrier benchmark

`shmem-barrier`: barrier benchmark

Usage: `shmem-barrier [options]`

Options: See Table 6-7

Table 6-7. QLogic SHMEM barrier benchmark options

Option	Description
`-h`	displays the help page
`-i INTEGER[K]`	outer iterations (default 1)

QLogic SHMEM reduce benchmark

`shmem-reduce`: reduce benchmark

Usage: shmem-reduce [options]

Options: See Table 6-8

Table 6-8. QLogic SHMEM reduce benchmark options

Option	Description
`-b INTEGER`	number of barriers between reduces (default 0)
`-h`	displays the help page
`-i INTEGER[K]`	outer iterations (default 1)
`-r INTEGER`	inner iterations (default 10000)

7 Virtual Fabric support in PSM

Introduction

Performance Scaled Messaging (PSM) provides support for full Virtual Fabric (vFabric) integration, allowing users to specify IB Service Level (SL) and Partition Key (PKey), or to provide a configured Service ID (SID) to target a vFabric. Support for using IB path record queries to the QLogic Fabric Manager during connection setup is also available, enabling alternative switch topologies such as Mesh/Torus. Note that this relies on the Distributed SA cache from FastFabric.

All PSM enabled MPIs can leverage these capabilities transparently, but only one MPI (Open MPI) is configured to support it natively. Native support here means that MPI specific mpirun switches are available to activate/deactivate these features. Other MPIs will require use of environment variables to leverage these capabilities. With MPI applications, the environment variables need to be propagated across all nodes/processes and not just the node from where the job is submitted/run. The mechanisms to do this are MPI specific, but for two common MPIs the following may be helpful:

- **Open MPI**: Use `-x ENV_VAR=ENV_VAL` in the mpirun command line. Example:

  ```
  mpirun -np 2 -machinefile machinefile -x
  PSM_ENV_VAR=PSM_ENV_VAL prog prog_args
  ```

- **MVAPICH2**: Use `mpirun_rsh` to perform job launch. Do not use `mpiexec` or `mpirun`. Specify the environment variable and value in the `mpirun` command line before the program argument.
  Example:

  ```
  mpirun_rsh -np 2 -hostfile machinefile
  PSM_ENV_VAR=PSM_ENV_VAL prog prog_args
  ```

Some of the features available require appropriate versions of associated software and firmware for correct operation. These requirements are listed in the relevant sections.

Virtual Fabric Support

Virtual Fabric (vFabric) in PSM is supported with the QLogic Fabric Manager. The latest version of the QLogic Fabric Manager contains a sample `qlogic_fm.xml` file with pre-configured vFabrics for PSM. Sixteen unique Service IDs have been allocated for PSM enabled MPI vFabrics to ease their testing however any Service ID can be used. Refer to the *QLogic Fabric Manager User Guide* on how to configure vFabrics.

There are two ways to use vFabric with PSM. The “legacy” method requires the user to specify the appropriate SL and Pkey for the vFabric in question. For complete integration with vFabrics, users can now specify a Service ID (SID) that identifies the vFabric to be used. PSM will automatically obtain the SL and Pkey to use for the vFabric from the QLogic Fabric Manager via path record queries.

Using SL and PKeys

SL and Pkeys can be specified natively for Open MPI. For other MPIs use the following list of environment variables to specify the SL and Pkey. The environment variables need to be propagated across all processes for correct operation.

NOTE

This is available with Open MPI v1.3.4rc4 and above only!

- **Open MPI**: Use mca parameters (`mtl_psm_ib_service_level` and `mtl_psm_ib_pkey`) to specify the pkey on the mpirun command line. Example:

  ```
  mpirun -np 2 -machinefile machinefile -mca
  mtl_psm_ib_service_level SL -mca mtl_psm_ib_pkey Pkey
  prog prog_args
  ```

- Other MPIs can use the following environment variables that are propagated across all processes. This process is MPI library specific but samples on how to do this for Open MPI and MVAPICH2 are listed in the “Introduction” on page 7-1.

  - `IPATH_SL=SL # Service Level to Use 0-15`
  - `PSM_PKEY=Pkey # Pkey to use`

Using Service ID

Full vFabric integration with PSM is available, allowing the user to specify a SID. For correct operation, PSM requires the following components to be available and configured correctly.

- QLogic host Fabric Manager Configuration – PSM MPI vFabrics need to be configured and enabled correctly in the `qlogic_fm.xml` file. 16 unique SIDs have been allocated in the sample file.
- OFED+ library needs to be installed on all nodes. This is available as part of Fast Fabrics tools.
- QLogic Distributed SA needs to be installed, configured and activated on all the nodes. This is part of FastFabrics tools. Please refer to *QLogic Fast Fabric User Guide* on how to configure and activate the Distributed SA. The SIDs configured in the QLogic Fabric Manager configuration file should also be provided to the Distributed SA for correct operation.

Service ID can be specified natively for Open MPI. For other MPIs use the following list of environment variables. The environment variables need to be propagated across all processes for correct operation.

- Open MPI: Use mca parameters (`mtl_psm_ib_service_id` and `mtl_psm_path_query`) to specify the service id on the mpirun command line. Example:

  ```
  mpirun -np 2 -machinefile machinefile -mca
  mtl_psm_path_query opp -mca mtl_psm_ib_service_id SID
  prog prog_args
  ```

- Other MPIs can use the following environment variables:

  - `PSM_PATH_REC=opp # Path record query mechanism to use. Always specify opp`
  - `PSM_IB_SERVICE_ID=SID # Service ID to use`

SL2VL mapping from the Fabric Manager

PSM is able to use the SL2VL table as programmed by the QLogic Fabric Manager. Prior releases required manual specification of the SL2VL mapping via an environment variable.

Verifying SL2VL tables on QLogic 7300 Series Adapters

`iba_saquery` can be used to get the SL2VL mapping for any given port however, QLogic 7300 series adapters exports the SL2VL mapping via sysfs files. These files are used by PSM to implement the SL2VL tables automatically. The SL2VL tables are per port and available under `/sys/class/infiniband/`*`hca name`*`/ports/`*`port #`*`/sl2vl`. The directory contains 16 files numbered 0-15 that specify the SL. Listing the SL files returns the VL as programmed by the SL.

8 Dispersive Routing

Infiniband® uses deterministic routing that is keyed from the Destination LID (DLID) of a port. The Fabric Manager programs the forwarding tables in a switch to determine the egress port a packet takes based on the DLID.

Deterministic routing can create hotspots even in full bisection bandwidth (FBB) fabrics for certain communication patterns if the communicating node pairs map onto a common upstream link, based on the forwarding tables. Since routing is based on DLIDs, the IB fabric provides the ability to assign multiple LIDs to a physical port using a feature called Lid Mask Control (LMC). The total number of DLIDs assigned to a physical port is 2^LMC with the LIDS being assigned in a sequential manner. The common IB fabric uses a LMC of 0, meaning each port has 1 LID assigned to it. With non-zero LMC fabrics, this results in multiple potential paths through the fabric to reach the same physical port. For example, multiple DLID entries in the port forwarding table that could map to different egress ports.

Dispersive routing, as implemented in the PSM, attempts to avoid congestion hotspots described above by “spraying” messages across these paths. A congested path will not bottleneck messages flowing down the alternate paths that are not congested. The current implementation of PSM supports fabrics with a maximum LMC of 3 (8 LIDs assigned per port). This can result in a maximum of 64 possible paths between a SLID, DLID pair ([SLID, DLID],[SLID, DLID+1], [SLID,DLID+2]…..[SLID,DLID+8],[SLID+1, DLID],[SLID+1, DLID+1]…..[SLID+7, DLID+8]). Keeping state associated with these many paths requires large amount of memory resources, with empirical data showing not much gain in performance beyond utilizing a small set of multiple paths. Therefore PSM reduces the number of paths actually used in the above case to 8 where the following paths are the only ones considered for transmission — [SLID, DLID], [SLID + 1, DLID + 1], [SLID + 2, DLID + 2] ….. [SLID + N, DLID + N]. This makes the resource requirements manageable while providing most of the benefits of dispersive routing (congestion avoidance by utilizing multiple paths).

Internally, PSM utilizes dispersive routing differently for small and large messages. Large messages are any messages greater-than or equal-to 64K. For large messages, the message is split into message fragments of 128K by default (called a window). Each of these message windows is sprayed across a distinct path between ports. All packets belonging to a window utilize the same path however the windows themselves can take a different path through the fabric. PSM assembles the windows that make up an MPI message before delivering it to the application. This allows limited out of order semantics through the fabrics to be maintain with little overhead. Small messages on the other hand always utilize a single path when communicating to a remote node however different processes executing on a node can utilize different paths for their communication between the nodes. For example, two nodes A and B each with 8 processors per node. Assuming the fabric is configured for a LMC of 3, PSM constructs 8 paths through the fabric as described above and a 16 process MPI application that spans these nodes (8 process per node). Then:

- Each MPI process is automatically bound to a given CPU core numbered between 0-7. PSM does this at startup to get improved cache hit rates and other benefits.
- Small Messages sent from a process on core N will use path N.

NOTE

Only path N will be used by this process for all communications to any process on the remote node.

- For a large message, each process will utilize all of the 8 paths and spray the windowed messages across it.

The above highlights the default path selection policy that is active in PSM when running on non-zero LMC configured fabrics. There are 3 other path selection policies that determine how to select the path (or path index from the set of available paths) used by a process when communicating with a remote node. The above path policy is called **adaptive**. The 3 remaining path policies are static policies that assign a static path on job startup for both small and large message transfers.

- **Static_Src**: Only **one path per process** is used for all remote communications. The path index is based on the CPU number the process is running.

NOTE

Multiple paths are still used in the fabric if multiple processes (each on a different CPU) are communicating.

- **Static_Dest**: The path selection is based on the CPU index of the destination process. Multiple paths can be used if data transfer is to different remote processes within a node. If multiple processes from Node A send a message to a single process on Node B only one path will be used across all processes.
- **Static_Base**: The only path that is used is the base path [SLID,DLID] between nodes regardless of the LMC of the fabric or the number of paths available. This is similar to how PSM operated till the IFS 5.1 release.

NOTE

A fabric configured with LMC of 0 even with the default adaptive policy enabled operates as the Static_Base policy as there only exists a single path between any pairs of port.

9 gPXE

gPXE is an open source (GPL) network bootloader. It provides a direct replacement for proprietary PXE ROMs. See http://etherboot.org/wiki/index.php for documentation and general information.

gPXE Setup

At least two machines and a switch are needed (or connect the two machines back-to-back and run QLogic Fabric Manager on the server).

- A DHCP server
- A boot server or http server (can be the same as the DHCP server)
- A node to be booted

  Use a QLE7340 or QLE7342 adapter for the node.

The following software is included with the QLogic OFED+ installation software package:

- gPXE boot image
- patch for DHCP server
- tool to install gPXE boot image in EPROM of card
- sample gPXE script

Everything that can be done with the proprietary PXE loader over Ethernet, can be done with the gPXE loader over IB. The gPXE boot code is only a mechanism to load an initial boot image onto the system. It is up to the downloaded boot image to do the rest.

For example, the boot image could be:

- A stand-alone memory test program
- A diskless kernel image that mounts its file systems via NFS
  Refer to http://www.faqs.org/docs/Linux-HOWTO/Diskless-HOWTO.html

- A Linux install image like kickstart, which then installs software to the local hard drive(s). Refer to http://www.faqs.org/docs/Linux-HOWTO/KickStart-HOWTO.html
- A second stage boot loader
- A live CD Linux image
- A gPXE script

Required Steps

1. Download a copy of the gPXE image.

   Located at:

   - The executable to flash the EXPROM on the QLogic IB adapters is located at: `/usr/sbin/ipath_exprom`
   - The gPXE driver for QLE7300 series IB adapters (the EXPROM image) is located at: `/usr/share/infinipath/gPXE/iba7322.rom`

2. In order for dhcpd to correctly load, assign IP addresses to the IB adapter GUID. The dhcpd on the existing DHCP server may need to be patched. This patch will be provided via the gPXE rpm installation.

3. Write the ROM image to the IB adapter.

   This only needs to be done once per IB adapter.

   ```
   ipath_exprom -e -w iba7xxx.rom
   ```

   In some cases, executing the above command results in a hang. If you experience a hang, type CTRL+C to quit, then execute one flag at a time:

   ```
   ipath_exprom -e iba7xxx.rom
   ipath_exprom -w iba7xxx.rom
   ```

4. Enable booting from the IB adapter (gPXE device) in the BIOS

Preparing the DHCP Server in Linux

When the boot session starts, the gPXE firmware attempts to bring up an adapter network link. If it succeeds to bring up a connected link, the gPXE firmware communicates with the DHCP server. The DHCP server assigns an IP address to the gPXE client and provides it with the location of the boot program.

Installing DHCP

gPXE requires that the DHCP server runs on a machine that supports IP over IB.

NOTE

Prior to installing DHCP, make sure that QLogic OFED+ is already installed on your DHCP server.

1. Download and install the latest DHCP server from www.isc.org.

   Standard DHCP fields holding MAC address are not large enough to contain an IPoIB hardware address. To overcome this problem, DHCP over IB messages convey a client identifier field used to identify the DHCP session. This client identifier field can be used to associate an IP address with a client identifier value, such that the DHCP server will grant the same IP address to any client that conveys this client identifier.

2. Unpack the latest downloaded DHCP server.

   ```
   tar zxf dhcp-release.tar.gz
   ```

3. Uncomment the line `/* #define USE_SOCKETS */` in `dhcp-release/includes/site.h`

4. Change to the main directory.

   ```
   cd dhcp-release
   ```

   NOTE

   If there is an older version of DHCP installed, save it before continuing with the following steps.

5. Configure the source.

   ```
   ./configure
   ```

6. When the configuration of DHCP is finished, build the DHCP server.

   ```
   make
   ```

7. When the DHCP has successfully finished building, install DHCP.

   ```
   make install
   ```

Configuring DHCP

1. From the client host, find the GUID of the HCA by using `p1info` or look at the GUID label on the IB adapter.

2. Turn the GUID into a MAC address and specify the port of the IB adapter that is going to be used at the end, using `b0` for port0 or `b1` for port1.

   For example for a GUID that reads `0x00117500005a6eec`, the MAC address would read: `00:11:75:00:00:5a:6e:ec:b0`

3. Add the MAC address to the DHCP server.

   The following is the sample `/etc/dhcpd.conf` file that specifies the HCA GUID for the hardware address:

```
#
# DHCP Server Configuration file.
#   see /usr/share/doc/dhcp*/dhcpd.conf.sample
#

ddns-update-style none;

subnet 10.252.252.0 netmask 255.255.255.0 {
     option subnet-mask 255.255.255.0;
     range dynamic-bootp 10.252.252.100 10.252.252.109;

     host hl5-0 {
          hardware unknown-32 00:11:75:00:00:7e:c1:b0;
          option host-name "hl5";
     }

     host hl5-1 {
          hardware unknown-32 00:11:75:00:00:7e:c1:b1;
          option host-name "hl5";
     }

     filename
     "http://10.252.252.1/images/uniboot/uniboot.php";
}
```

   In this example, host `hl5` has a dual port IB adapter. hl5-0 corresponds to port 0, and `hl5-1` corresponds to port 1 on the adapter.

4. Restart the DHCP server

Netbooting Over IB

The following procedures are an example of netbooting over IB, using an HTTP boot server.

Prerequisites

- Required steps from above have been executed.
- The BIOS has been configured to enable booting from the IB adapter. The gPXE IB device should be listed as the first boot device.
- Apache server has been configured with PHP on your network, and is configured to serve pages out of `/vault`.
- It is understood in this example that users would have their own tools and files for diskless booting with an http boot server.

NOTE

The dhcpd and apache configuration files referenced in this example are included as examples, and are not part of the QLogic OFED+ installed software. Your site boot servers may be different, see their documentation for equivalent information.

Instructions on installing and configuring a dhcp server or a boot server are beyond the scope of this document.

Boot Server Setup

Configure the boot server for your site.

NOTE

gPXE supports several file transfer methods such as TFTP, HTTP, iSCSI. This example uses HTTP since it generally scales better and is the preferred choice.

NOTE

This step involves setting up a http server and needs to be done by a user that understands server setup on the http server is being used

1. Install Apache.

2. Create an `images.conf` file and a `kernels.conf` file and place them in the `/etc/httpd/conf.d` directory. This sets up aliases for and tells apache where to find them:

   ```
   /images – http://10.252.252.1/images/
   /kernels – http://10.252.252.1/kernels/
   ```

   The following is an example of the `images.conf` file

   ```
   Alias /images /vault/images

   <Directory "/vault/images">
       AllowOverride All
       Options Indexes FollowSymLinks
       Order allow,deny
       Allow from all
   </Directory>
   ```

   The following is an example of the `kernels.conf` file

   ```
   Alias /kernels /boot

   <Directory "/boot">
       AllowOverride None
       Order allow,deny
       Allow from all
   </Directory>
   ```

3. Make a `uniboot` directory:

   ```
   mkdir -p /vault/images/uniboot
   ```

4. Create a `initrd.img file`

   Prerequisites

   - “gPXE Setup” on page 9-1 has been completed.
   - “Preparing the DHCP Server in Linux” on page 9-2 has been completed

To add an IB driver into the `initrd` file, The IB modules need to be copied to the diskless image. The host machine needs to be pre-installed with the QLogic OFED+ Host Software that is appropriate for the kernel version the diskless image will run. The QLogic OFED+ Host Software is available for download from http://driverdownloads.qlogic.com/QLogicDriverDownloads_UI/default.aspx

NOTE

The remainder of this section assumes that QLogic OFED+ has been installed on the Host machine.

⚠WARNING

The following procedure modifies critical files used in the boot procedure. It must be executed by users with expertise in the boot process. Improper application of this procedure may prevent the diskless machine from booting.

a. If `/vault/images/initrd.img` file is already present on the server machine, back it up. For example:

```
cp -a /vault/images/initrd.img /vault/images/
initrd.img.bak
```

b. The infinipath rpm will install the file `/usr/share/infinipath/gPXE/gpxe-qib-modify-initrd` with contents similar to the following example. You can either run the script to generate a new initrd image, or use it as an example, and customize as appropriate for your site.

```
# This assumes you will use the currently running version
of linux, and
# that you are starting from a fully configured machine of
the same type
# (hardware configuration), and BIOS settings.
#
# start with a known path, to get the system commands
PATH=/sbin:/usr/sbin:/bin:/usr/bin:$PATH

# start from a copy of the current initd image

mkdir -p /var/tmp/initrd-ib
cd /var/tmp/initrd-ib

kern=$(uname -r)
if [ -e /boot/initrd-${kern}.img ]; then
initrd=/boot/initrd-${kern}.img
elif [ -e /boot/initrd ]; then
initrd=/boot/initrd
else
echo Unable to locate correct initrd, fix script and
re-run
exit 1
fi

cp ${initrd} initrd-ib-${kern}.img

# Get full original listing
gunzip -dc initrd-ib-${kern}.img | cpio -it --quiet |
grep -v '^\.$' | sort -o Orig-listing

# start building modified image
rm -rf new # for retries
mkdir new
cd new
```

```
# extract previous contents
gunzip -dc ../initrd-ib-${kern}.img | cpio --quiet -id

# add infiniband modules
mkdir -p lib/ib
find /lib/modules/${kern}/updates -type f | \
 egrep
'(iw_cm|ib_(mad|addr|core|sa|cm|uverbs|ucm|umad|ipoib|qib
).ko|rdma_|ipoib_helper)' | \
 xargs -I '{}' cp -a  '{}' lib/ib

# Some distros have ipoib_helper, others don't require it
if [ -e lib/ib/ipoib_helper ]; then
helper_cmd='/sbin/insmod /lib/ib/ipoib_helper.ko'
fi

# On some kernels, the qib driver will require the dca
module
if modinfo -F depends ib_qib | grep -q dca; then
  cp $(find /lib/modules/$(uname -r) -name dca.ko) lib/ib
  dcacmd='/sbin/insmod /lib/ib/dca.ko'
else
  dcacmd=
fi

# IB requires loading an IPv6 module. If you do not have
it in your initrd, add it
if grep -q ipv6 ../Orig-listing; then
# already added, and presumably insmod'ed, along with any
dependencies
v6cmd=
else
echo -e 'Adding IPv6 and related modules\n'
cp /lib/modules/${kern}/kernel/net/ipv6/ipv6.ko lib
IFS=' ' v6cmd='echo "Loading IPV6"
/sbin/insmod /lib/ipv6.ko'
# Some versions of IPv6 have dependencies, add them.
xfrm=$(modinfo -F depends ipv6)
if [ ${xfrm} ]; then
    cp $(find /lib/modules/$(uname -r) -name ${xfrm}.ko)
lib
```

```
    IFS=' ' v6cmd='/sbin/insmod /lib/'${xfrm}'.ko
'"$v6cmd"
    crypto=$(modinfo -F depends $xfrm)
    if [ ${crypto} ]; then
      cp $(find /lib/modules/$(uname -r) -name
${crypto}.ko) lib
    IFS=' ' v6cmd='/sbin/insmod /lib/'${crypto}'.ko
'"$v6cmd"
    fi
fi
fi

# we need insmod to load the modules; if not present it,
copy it
mkdir -p sbin
grep -q insmod ../Orig-listing || cp /sbin/insmod sbin

echo -e 'NOTE: you will need to config ib0 in the normal
way in your booted root
filesystem, in order to use it for NFS, etc.\n'

# Now build the commands to load the additional modules.
We add them just after
# the last existing insmod command, so all other
dependences will be resolved
# You can change the location if desired or necessary.
# loading order is important.  You can verify the order
works ahead of time
# by running "/etc/init.d/openibd stop", and then running
these commands
# manually by cut and paste
# This will work on SLES, although different than the
standard mechanism

cat > ../init-cmds << EOF
# Start of IB module block
$v6cmd
echo "loading IB modules"
/sbin/insmod /lib/ib/ib_addr.ko
/sbin/insmod /lib/ib/ib_core.ko
/sbin/insmod /lib/ib/ib_mad.ko
```

```
/sbin/insmod /lib/ib/ib_sa.ko
/sbin/insmod /lib/ib/ib_cm.ko
/sbin/insmod /lib/ib/ib_uverbs.ko
/sbin/insmod /lib/ib/ib_ucm.ko
/sbin/insmod /lib/ib/ib_umad.ko
/sbin/insmod /lib/ib/iw_cm.ko
/sbin/insmod /lib/ib/rdma_cm.ko
/sbin/insmod /lib/ib/rdma_ucm.ko
$dcacmd
/sbin/insmod /lib/ib/ib_qib.ko
$helper_cmd
/sbin/insmod /lib/ib/ib_ipoib.ko
echo "finished loading IB modules"
# End of IB module block
EOF

# first get line number where we append (after last insmod
if any, otherwse
# at start
line=$(egrep -n insmod init | sed -n '$s/:.*//p')
if [ ! "${line}" ]; then line=1; fi
sed -e "${line}r ../init-cmds" init > init.new



# show the difference, then rename
echo -e 'Differences between original and new init
command script\n'
diff init init.new
mv init.new init
chmod 700 init

# now rebuilt the initrd image
find . | cpio --quiet -H newc -o | gzip >
../initrd-${kern}.img

cd ..

# get the file list in the new image
gunzip -dc initrd-${kern}.img | cpio --quiet -it | grep
-v '^\.$' | sort -o New-listing
```

```
# and show the differences.
echo -e '\nChanges in files in initrd image\n'
diff Orig-listing New-listing

# copy the new initrd to wherever you have configure the
dhcp server to look
# for it (here we assume it's /images)
mkdir -p /images
cp  initrd-${kern}.img /images

echo -e '\nCompleted initrd for IB'
ls -l /images/initrd-${kern}.img
```

c. Run the `usr/share/infinipath/gPXE/gpxe-qib-modify-initrd` script to create the `initrd.img` file.

At this stage, the `initrd.img` file is ready and located at the location where the DHCP server was configured to look for it.

5. Create a `uniboot.php` file and save it to `/vault/images/uniboot.`

NOTE

The `uniboot.php` generates a gPXE script that will attempt to boot from the `/boot/vmlinuz-2.6.18-128.el5` kernel. If you want to boot from a different kernel, edit `uniboot.php` with the appropriate kernel string in the `$kver` variable.

The following is an example of a `uniboot.php` file:

```
<?
header ( 'Content-type: text/plain' );

function strleft ( $s1, $s2 ) {
  return substr ( $s1, 0, strpos ( $s1, $s2 ) );
}

function baseURL() {
  $s = empty ( $_SERVER["HTTPS"] ) ? '' :
    ( $_SERVER["HTTPS"] == "on" ) ? "s" : "";
  $protocol = strleft ( strtolower (
$_SERVER["SERVER_PROTOCOL"] ), "/" ).$s;
  $port = ( $_SERVER["SERVER_PORT"] == "80" ) ? "" :
    ( ":".$_SERVER["SERVER_PORT"] );
  return $protocol."://".$_SERVER['SERVER_NAME'].$port;
}

$baseurl = baseURL();
$selfurl = $baseurl.$_SERVER['REQUEST_URI'];
$dirurl = $baseurl.( dirname ( $_SERVER['SCRIPT_NAME'] ) );

$kver = "2.6.18-164.11.1.el5";

echo <<< EOF
#!gpxe
initrd /images/initrd.img
kernel /kernels/vmlinuz-${kver} bootfile=${selfurl}
ip=\${net0/ip}::\${net0/gateway}:\${net0/netmask}:\${net0/hos
tname}:ib0:off vga=788 console=tty0 console=ttyS0,115200
debug root=/dev/hdb1
boot
EOF;
?>
```

The generated gPXE script tells gPXE to load `/boot/vmlinuz-2.6.18-128.el5` and `/vault/images/initrd.img` files from the httpd server node and run them.

1. Copy `vmlinuz-2.6.18-128.el5` to `/boot` on the boot server.

This is the kernel that will boot.

This file can be copied from any machine that has RHEL5.3 installed.

2. Start httpd

Steps on the gPXE Client

1. Ensure that the HCA is listed as the first bootable device in the BIOS.
2. Reboot the test node(s) and enter the BIOS boot setup.

   This is highly dependent on the BIOS for the system but you should see a menu for boot options and a submenu for boot devices.

   Select gPXE IB as the first boot device.

   When you power on the system or press the reset button, the system will execute the boot code on the HCA that will query the DHCP server for the IP address and boot image to download.

   Once the boot image is downloaded, the BIOS/HCA is finished and the boot image is ready.
3. Verify system boots off of the kernel image on the boot server. The best way to do this is to boot into a different kernel from the one installed on the hard drive on the client, or to un-plug the hard drive on the client and verify that on boot up, a kernel and file system exist.

HTTP Boot Setup

gPXE supports booting diskless machines. To enable using an IB driver, the (remote) kernel or initrd image must include and be configured to load that driver. This can be achieved either by compiling the HCA driver into the kernel, or by adding the device driver module into the initrd image and loading it.

1. Make a new directory

   ```
   mdir /vault/images/uniboot
   ```
2. Change directories

   ```
   cd /vault/images/uniboot
   ```
3. Create a `initrd.img file` using the information and example in Step 4 of Boot Server Setup.
4. Create a `uniboot.php` file using the example in Step 4 of Boot Server Setup.

5. Create an `images.conf` file and a `kernels.conf` file using the examples in Step 2 of Boot Server Setup and place them in the `/etc/httpd/conf.d` directory.

6. Edit `/etc/dhcpd.conf` file to boot the clients using HTTP

```
filename "http://172.26.32.9/images/uniboot/uniboot.php";
```

7. Restart the DHCP server

8. Start HTTP if it is not already running:

```
/etc/init.d/httpd start
```

A Benchmark Programs

Several MPI performance measurement programs are installed by default with the MPIs you choose to install (such as Open MPI, MVAPICH2 or MVAPICH). This appendix describes a few of these benchmarks and how to run them. Several of these programs are based on code from the group of Dr. Dhabaleswar K. Panda at the Network-Based Computing Laboratory at the Ohio State University. For more information, see: http://mvapich.cse.ohio-state.edu/

These programs allow you to measure MPI latency, bandwidth, and message rate between two or more nodes in your cluster. The executables are installed by default under `/usr/mpi/compiler/mpi/tests/osu_benchmarks-3.1.1`. The remainder of this chapter will assume that the gcc-compiled version of Open MPI was installed in the default location of `/usr/mpi/gcc/openmpi-1.4.3-qlc` and that mpi-selector is used to choose this Open MPI version as the MPI to be used.

The following examples are intended to show only the syntax for invoking these programs and the meaning of the output. They are not representations of actual IB performance characteristics.

For additional MPI sample applications refer to Section 5 of the *QLogic FastFabric Command Line Interface Reference Guide*.

Benchmark 1: Measuring MPI Latency Between Two Nodes

In the MPI community, *latency for a message of given size* is the time difference between a node program's calling `MPI_Send` and the time that the corresponding `MPI_Recv` in the receiving node program returns. The term *latency*, alone without a qualifying message size, indicates the latency for a message of size zero. This latency represents the minimum overhead for sending messages, due to both software overhead and delays in the electronics of the fabric. To simplify the timing measurement, latencies are usually measured with a *ping-pong* method, timing a round-trip and dividing by two.

The program `osu_latency`, from Ohio State University, measures the latency for a range of messages sizes from 0bytes to 4 megabytes. It uses a ping-pong method, where the rank zero process initiates a series of sends and the rank one process echoes them back, using the blocking MPI send and receive calls for all operations. Half the time interval observed by the rank zero process for each exchange is a measure of the latency for messages of that size, as previously defined. The program uses a loop, executing many such exchanges for each message size, to get an average. The program defers the timing until the message has been sent and received a number of times, to be sure that all the caches in the pipeline have been filled.

This benchmark always involves two node programs. It can be run with the command:

```
$ mpirun -H host1,host2 \
/usr/mpi/gcc/openmpi-1.4.3-qlc/tests/osu_benchmarks-3.1.1/osu
_latency
```

`-H` (or `--hosts`) allows the specification of the host list on the command line instead of using a host file (with the `-m` or `-machinefile` option). Since only two hosts are listed, this implies that two host programs will be started (as if `-np 2` were specified). The output of the program looks like:

```
# OSU MPI Latency Test v3.1.1)
# Size          Latency (us)
0               1.67
1               1.68
2               1.69
4               1.68
8               1.68
16              1.93
32              1.92
64              1.92
128             1.99
256             2.12
512             2.38
1024            2.74
2048            3.52
4096            4.59
8192            6.52
16384           9.98
32768           17.65
65536           52.11
131072          84.07
262144          114.90
524288          241.97
1048576         422.41
2097152         783.21
4194304         1596.37
```

The first column displays the message size in bytes. The second column displays the average (one-way) latency in microseconds. This example shows the syntax of the command and the format of the output, and is not meant to represent actual values that might be obtained on any particular IB installation.

Benchmark 2: Measuring MPI Bandwidth Between Two Nodes

The `osu_bw` benchmark measures the maximum rate that you can pump data between two nodes. This benchmark also uses a ping-pong mechanism, similar to the `osu_latency` code, except in this case, the originator of the messages pumps a number of them (64 in the installed version) in succession using the non-blocking MPI_I send function, while the receiving node consumes them as quickly as it can using the non-blocking MPI_Irecv function, and then returns a zero-length acknowledgement when all of the sent data has been received. You can run this program by typing:

```
$ mpirun -H host1,host2 \

/usr/mpi/gcc/openmpi-1.4.3-qlc/tests/osu_benchmarks-3.1.1/osu
_bw
```

Typical output might look like:

```
# OSU MPI Bandwidth Test v3.1.1
# Size        Bandwidth (MB/s)
1                         2.35
2                         4.69
4                         9.38
8                        18.80
16                       34.55
32                       68.89
64                      137.87
128                     265.80
256                     480.19
512                     843.70
1024                   1353.48
2048                   1984.11
4096                   2152.61
8192                   2249.00
16384                  2680.75
32768                  2905.83
65536                  3170.05
131072                 3224.15
262144                 3241.35
524288                 3270.21
1048576                3286.05
2097152                3292.64
4194304                3283.87
```

The increase in measured bandwidth with the messages' size is because the contribution of each packet's overhead to the measured time becomes relatively smaller.

Benchmark 3: Messaging Rate Microbenchmarks

OSU Multiple Bandwidth / Message Rate test (osu_mbw_mr)

`osu_mbw_mr` is a multi-pair bandwidth and message rate test that evaluates the aggregate uni-directional bandwidth and message rate between multiple pairs of processes. Each of the sending processes sends a fixed number of messages (the window size) back-to-back to the paired receiving process before waiting for a reply from the receiver. This process is repeated for several iterations. The objective of this benchmark is to determine the achieved bandwidth and message rate from one node to another node with a configurable number of processes running on each node. You can run this program as follows:

```
$ mpirun -H host1,host2 -npernode 12 \
/usr/mpi/gcc/openmpi-1.4.3-qlc/tests/osu_benchmarks-3.1.1/osu
_mbw_mr
```

This was run on 12-core compute nodes, so we used Open MPI's `-npernode 12` option to place 12 MPI processes on each node (for a total of 24) to maximize message rate. Note that the output below indicates that there are 12 pairs of communicating processes.

```
# OSU MPI Multiple Bandwidth / Message Rate Test v3.1.1
# [ pairs: 12 ] [ window size: 64 ]
# Size                     MB/s          Messages/s
1                         22.77         22768062.43
2                         44.90         22449128.66
4                         91.75         22938300.02
8                        179.23         22403849.44
16                       279.91         17494300.07
32                       554.16         17317485.47
64                      1119.88         17498101.32
128                     1740.54         13597979.96
256                     2110.22          8243066.36
512                     2353.17          4596038.46
1024                    2495.88          2437386.38
2048                    2573.99          1256833.08
4096                    2567.88           626923.21
8192                    2757.54           336613.42
16384                   3283.94           200435.90
32768                   3291.54           100449.84
65536                   3298.20            50326.50
131072                  3305.77            25221.05
262144                  3310.39            12628.14
524288                  3310.83             6314.90
1048576                 3311.11             3157.72
2097152                 3323.50             1584.77
4194304                 3302.35              787.34
```

An Enhanced Multiple Bandwidth / Message Rate test (mpi_multibw)

`mpi_multibw` is a version of osu_mbw_mr which has been enhanced by QLogic to, optionally, run in a bidirectional mode and to scale better on the larger multi-core nodes available today This benchmark is a modified form of the OSU Network-Based Computing Lab's `osu_mbw_mr` benchmark (as shown in the previous example). It has been enhanced with the following additional functionality:

- N/2 is dynamically calculated at the end of the run.
- You can use the `-b` option to get a bidirectional message rate and bandwidth results.
- Scalability has been improved for larger core-count nodes.

The benchmark has been updated with code to dynamically determine what processes are on which host. Thefollowing is an example output when running mpi_multibw:

```
$ mpirun -H host1,host2 -npernode 12  \
  /usr/mpi/gcc/openmpi-1.4.3-qlc/tests/qlogic/mpi_multibw
# PathScale Modified OSU MPI Bandwidth Test
(OSU Version 2.2, PathScale $Revision: 1.1.2.1 $)
# Running on 12 procs per node (uni-directional traffic for
each process pair)

# Size          Aggregate Bandwidth (MB/s)      Messages/s
1               24.992623                       24992622.996615
2               50.015847                       25007923.312888
4               100.075479                      25018869.818990
8               200.115037                      25014379.610716
16              284.475601                      17779725.040265
32              568.950239                      17779694.953511
64              1137.899392                     17779677.998115
128             1758.183987                     13735812.394705
256              2116.159352                     8266247.468294
512              2355.027827                     4599663.724469
1024             2496.960650                     2438438.134886
2048             2574.260975                     1256963.366877
4096             2567.861960                      626919.423819
8192             2746.514440                      335267.875961
16384            3284.264487                      200455.596122
32768            3292.007839                      100464.106405
65536            3299.800622                      50350.961641
131072           3306.998105                      25230.393259
262144           3309.840069                      12626.037860
524288           3323.339300                      6338.766671
1048576          3323.068802                      3169.125368
2097152          3307.077899                      1576.937627
4194304          3300.327382                      786.859365
Searching for N/2 bandwidth.  Maximum Bandwidth of 3323.339300
MB/s...
Found N/2 bandwidth of 1662.009095 MB/s at size 121 bytes
```

Note the improved message rate at small message sizes of ~25 million compared to the rate of 22.8 million measured with `osu_mbw_mr`. Also note that it only takes a message of size 121 bytes to generate half of the peak uni-directional bandwidth.

The following is an example output when running with the bidirectional option (`-b`):

```
$ mpirun -H host1,host2 -np 24  \
  /usr/mpi/gcc/openmpi-1.4.3-qlc/tests/qlogic/mpi_multibw -b
# PathScale Modified OSU MPI Bandwidth Test
(OSU Version 2.2, PathScale $Revision: 1.1.2.1 $)
# Running on 12 procs per node (bi-directional traffic for
each process pair)
# Size          Aggregate Bandwidth (MB/s)      Messages/s
1               34.572819                       34572819.324348
2               68.984920                       34492459.942272
4               137.870850                      34467712.532016
8               274.914966                      34364370.730843
16              438.182185                      27386386.585309
32              871.077525                      27221172.671073
64              1743.576039                     27243375.616870
128             3046.774606                     23802926.607917
256             3968.178042                     15500695.477711
512             4558.456908                      8903236.148204
1024            4876.777738                      4762478.259397
2048            5050.255245                      2465944.943769
4096            5063.142612                      1236118.801851
8192            5234.475557                       638974.066993
16384           6255.483598                       381804.418801
32768           6236.354159                       190318.425252
65536           6288.370045                        95952.912066
131072          6330.494823                        48297.842586
262144          6351.690777                        24229.777437
524288          6353.021307                        12117.426504
1048576         6353.890433                         6059.542115
2097152         6353.951840                         3029.800339
4194304         6354.671923                         1515.071851
Searching for N/2 bandwidth.  Maximum Bandwidth of 6354.671923
MB/s...
Found N/2 bandwidth of 3184.322181 MB/s at size 170 bytes
```

The higher peak bi-directional messaging rate of 34.6 million messages per second at the 1 byte size, compared to 25 million messages/sec. when run unidirectionally.

B SRP Configuration

SRP Configuration Overview

SRP stands for SCSI RDMA Protocol. It allows the SCSI protocol to run over IB for Storage Area Network (SAN) usage. SRP interfaces directly to the Linux file system through the SRP Upper Layer Protocol (ULP). SRP storage can be treated as another device.

In this release, two versions of SRP are available: QLogic SRP and OFED SRP. QLogic SRP is available as part of the QLogic OFED Host Software, QLogic IFS, Rocks Roll, and Platform PCM downloads.

SRP has been tested on targets from DataDirect Networks™ and Engenio™ (now LSI Logic®).

NOTE

Before using SRP, the SRP targets must already be set up by your system administrator.

Important Concepts

- A SRP Initiator Port is a adapter port through which the host communicates with a SRP target device (e.g., a Fibre Channel disk array) via a SRP target port.
- A SRP Target Port is an IOC of the VIO hardware. In the context of VIO hardware, an IOC can be thought of as a SRP target. An FVIC contains 2 IOCs. IOC1 maps to the first adapter on the FVIC, and IOC2 maps to the 2nd adapter on the FVIC. On an FCBM, there are also 2 IOCs, and IOC1 maps to port 1 of the adapter of the FC BM and IOC2 maps to port 2 of the adapter of the FC BM.
- A Fibre Channel Target Device is a device containing storage resources that is located remotely from a Fibre Channel host. In the context of SRP/VIO hardware, this is typically an array of disks connected via Fibre Channel to the VIO hardware.

- A SRP Initiator Extension is a 64-bit numeric value that is appended to the port GUID of the SRP initiator port, which allows an SRP initiator port to have multiple SRP maps associated with it. Maps are for FVIC only. IB attached storage will use their own mechanism as maps are not necessary.
- A SRP Initiator is the combination of an SRP initiator port and an SRP initiator extension.
- A SRP Target is identified by the combination of an SRP target IOC and a SRP target extension.
- A SRP Session defines a connection between an SRP initiator and a SRP target.
- A SRP Map associates an SRP session with a Fibre Channel Target Device. This mapping is configured on the VIO hardware. Maps are for FVIC only. IB attached storage will use their own mechanism as maps are not necessary.

NOTE

- If a device connected to a map is changed, the SRP driver must be restarted.
- If the connected device is unreachable for a period of time, the Linux kernel may set the device offline. If this occurs the SRP driver must be restarted.

- A SRP Adapter is a collection of SRP sessions. This collection is then presented to the Linux kernel as if those sessions were all from a single adapter. All sessions configured for an adapter must ultimately connect to the same target device.

NOTE

- The SRP driver must be stopped before OFED (i.e., `openibd`) is stopped or restarted. This is due to SRP having references on OFED modules. The Linux kernel will not allow those OFED modules to be unloaded.

QLogic SRP Configuration

The QLogic SRP is installed as part of the QLogic OFED+ Host Software or the QLogic IFS. The following section provide procedures to set up and configure the QLogic SRP.

Stopping, Starting and Restarting the SRP Driver

To stop the qlgc_srp driver, use the following command:

```
/etc/init.d/qlgc_srp stop
```

To start the qlgc_srp driver, use the following command:

```
/etc/init.d/qlgc_srp start
```

To restart the qlgc_srp driver, use the following command:

```
/etc/init.d/qlgc_srp restart
```

Specifying a Session

In the SRP configuration file, a session command is a block of configuration commands, surrounded by `begin` and `end` statements. Sessions can be specified in several different ways, but all consist of specifying an SRP initiator and an SRP target port. For example:

```
session
begin
card: 0
port: 1
targetIOCGuid: 0x00066AXXXXXXXXXX
initiatorExtension: 2
end
```

The session command has two parts; the part that specifies the SRP initiator and the part that specifies the SRP target port. The SRP initiator contains two parts, the SRP initiator port and the SRP initiator extension. The SRP initiator extension portion of the SRP initiator is optional, and defaults to a value of 1. However, if a SRP initiator extension is not specified, each port on the adapter can use only one SRP map per VIO device. In addition a `targetExtension` can be specified (the default is **`1`**).

The SRP Initiator Port may be specified in two different ways:

1. By using the port GUID of the adapter port used for the connection, or
2. Specify the index of the adapter card being used (this is zero-based, so if there is only one adapter card in the system use a value of `0`) and the index of the port number (1 or 2) of the adapter card being used.

The SRP target port may be specified in two different ways:

1. By the port GUID of the IOC, or

2. By the IOC profile string that is created by the VIO device (i.e., a string containing the chassis GUID, the slot number and the IOC number). FVIC creates the device in this manner, other devices have their own naming method.

To specify the host IB port to use, the user can either specify the port GUID of the local IB port, or simply use the index numbers of the cards and the ports on the cards. Cards are numbered from 0 on up, based on the order they occur in the PCI bus. Ports are numbered in the same way, from first to last. To see which cards and ports are available for use, type the following command:

```
ib_qlgc_srp_query
```

The system returns input similar to the following:

```
st187:~/qlgc-srp-1_3_0_0_1 # ib_qlgc_srp_query
QLogic Corporation. Virtual HBA (SRP) SCSI Query Application,
version 1.3.0.0.1

1 IB Host Channel Adapter present in system.
HCA Card 0      : 0x0002c9020026041c
Port 1 GUID     : 0x0002c9020026041d
Port 2 GUID     : 0x0002c9020026041e

SRP Targets     :
SRP IOC Profile : FVIC in Chassis 0x00066a000300012a, Slot 17,
Ioc 1
  SRP IOC GUID  : 0x00066a01dd000021
  SRP IU SIZE   : 320
  SRP IU SG SIZE: 15
  SRP IO CLASS  : 0xff00
  service 0     : name SRP.T10:0000000000000001 id
0x0000494353535250
  service 1     : name SRP.T10:0000000000000002 id
0x0000494353535250
  service 2     : name SRP.T10:0000000000000003 id
0x0000494353535250
  service 3     : name SRP.T10:0000000000000004 id
0x0000494353535250
  Target Path(s):
   HCA 0 Port 1 0x0002c9020026041d -> Target Port GID
0xfe80000000000000066a11dd000021
   HCA 0 Port 2 0x0002c9020026041e -> Target Port GID
0xfe80000000000000066a11dd000021

SRP IOC Profile : FVIC in Chassis 0x00066a000300012a, Slot 17,
Ioc 2
  SRP IOC GUID  : 0x00066a02dd000021
  SRP IU SIZE   : 320
  SRP IU SG SIZE: 15
  SRP IO CLASS  : 0xff00
  service 0     : name SRP.T10:0000000000000001 id
0x0000494353535250
  service 1     : name SRP.T10:0000000000000002 id
0x0000494353535250
  service 2     : name SRP.T10:0000000000000003 id
```

```
0x0000494353535250
  service 3     : name SRP.T10:0000000000000004 id
0x0000494353535250
  Target Path(s):
   HCA 0 Port 1 0x0002c9020026041d -> Target Port GID
0xfe80000000000000066a21dd000021
   HCA 0 Port 2 0x0002c9020026041e -> Target Port GID
0xfe80000000000000066a21dd000021

SRP IOC Profile : Chassis 0x00066A0050000135, Slot 5, IOC 1
  SRP IOC GUID  : 0x00066a013800016c
  SRP IU SIZE   : 320
  SRP IU SG SIZE: 15
  SRP IO CLASS  : 0xff00
  service 0     : name SRP.T10:0000000000000001 id
0x0000494353535250
  service 1     : name SRP.T10:0000000000000002 id
0x0000494353535250
  service 2     : name SRP.T10:0000000000000003 id
0x0000494353535250
  service 3     : name SRP.T10:0000000000000004 id
0x0000494353535250
  Target Path(s):
   HCA 0 Port 1 0x0002c9020026041d -> Target Port GID
0xfe80000000000000066a026000016c
   HCA 0 Port 2 0x0002c9020026041e -> Target Port GID
0xfe80000000000000066a026000016c
```

Determining the values to use for the configuration

In order to build the configuration file, use the command `ib_qlgc_srp_build_cfg` script as follows:

Enter `ib_qlgc_srp_build_cfg`. The system provides output similar to the following:

```
# qlgc_srp.cfg file generated by
/usr/sbin/ib_qlgc_srp_build_cfg, version 1.3.0.0.17, on Mon
Aug 25 13:42:16 EDT 2008
#Found QLogic OFED SRP
registerAdaptersInOrder: ON
#
================================================================
# IOC Name: BC2FC in Chassis 0x0000000000000000, Slot 6, Ioc 1
# IOC GUID: 0x00066a01e0000149 SRP IU SIZE : 320

#  service 0     : name SRP.T10:0000000000000001 id
0x0000494353535250
session
begin
  card: 0
  port: 1
  #portGuid: 0x0002c9030000110d
  initiatorExtension: 1
  targetIOCGuid: 0x00066a01e0000149
  targetIOCProfileIdString: "FVIC in Chassis
0x0000000000000000, Slot 6, Ioc 1"
  targetPortGid: 0xfe8000000000000000066a01e0000149
  targetExtension: 0x0000000000000001
  SID: 0x0000494353535250
  IOClass: 0xff00
end

adapter
begin
  adapterIODepth: 1000
  lunIODepth: 16
  adapterMaxIO: 128
  adapterMaxLUNs: 512
  adapterNoConnectTimeout: 60
  adapterDeviceRequestTimeout: 2
  # set to 1 if you want round robin load balancing
  roundrobinmode: 0
  # set to 1 if you do not want target connectivity
verification
```

```
  noverify: 0
  description: "SRP Virtual HBA 0"
end
```

The `ib_qlgc_srp_build_cfg` command creates a configuration file based on discovered target devices. By default, the information is sent to `stdout`. In order to create a configuration file, output should be redirected to a disk file. Enter `ib_qlgc_srp_build_cfg` **`-h`** for a list and description of the option flags.

NOTE

The default configuration generated by `ib_qlgc_srp_build_cfg` for OFED is similar to the one generated for the QuickSilver host stack with the following differences:

- For OFED, the configuration automatically includes `IOClass`
- For OFED, the configuration automatically includes `SID`
- For OFED, the configuration provides information on `targetPortGid` instead of `targetPortGuid`
- For OFED, the configuration automatically includes `targetIOCProfileIdString`.

Specifying an SRP Initiator Port of a Session by Card and Port Indexes

The following example specifies a session by card and port indexes. If the system contains only one adapter, use this method.

```
session
begin
#Specifies the near side by card index
card: 0 #Specifies first HCA
port: 1 #Specifies first port
targetIOCGuid: 0x00066013800016C
end
```

Specifying an SRP Initiator Port of Session by Port GUID

The following example specifies a session by port GUID. If the system contains more than one adapter, use this method.

```
session
begin
portGuid: 0x00066A00a00001a2 #Specifies port by its GUID
targetIOCGuid: 0x00066A013800016C
end
```

NOTE

When using this method, if the port GUIDs are changed, they must also be changed in the configuration file.

Specifying a SRP Target Port

The SRP target can be specified in two different ways. To connect to a particular SRP target no matter where it is in the fabric, use the first method (By IOCGUID). To connect to a SRP target that is in a certain chassis/slot, no matter which card it is on (For FVIC, the user does not want to change the configuration, if cards are switched in a slot) then use the second method.

1. By IOCGUID. For example:

```
targetIOCGuid: 0x00066A013800016c
```

2. By target IOC Profile String. For example:

```
targetIOCProfileIdString: "FVIC in Chassis
0x00066A005000010E, Slot 1, IOC 1"
```

NOTE

When specifying the `targetIOCProfileIdString`, the string is case and format sensitive. The easiest way to get the correct format is to cut and paste it from the output of the `/usr/sbin/ib_qlgc_srp_query` program.

NOTE

For FVIC, by specifying the SRP Target Port by IOCGUID, this ensures that the session will always be mapped to the specific port on this specific VIO hardware card, even if the card is moved to a different slot in the same chassis or even if it is moved to a different chassis.

NOTE

For FVIC, by specifying the SRP Target Port by Profile String, this ensures that the session will always be mapped to the VIO hardware card in the specific slot of a chassis, even if the VIO hardware card currently in that slot is replaced by a different VIO hardware card.

Specifying a SRP Target Port of a Session by IOCGUID

The following example specifies a target by IOC GUID:

```
session
begin
card: 0
port: 1
targetIOCGuid: 0x00066A013800016c #IOC GUID of the InfiniFibre
port
end
```

- 0x00066a**1**0dd000046
- 0x00066a**2**0dd000046

Specifying a SRP Target Port of a Session by Profile String

The following example specifies a target by Profile String:

```
session
begin
card: 0
port: 1
# FVIC in Chassis 0x00066A005000010E,
# Slot number 1, port 1
targetIOCProfileIdString: "FVIC in Chassis
0x00066A005000010E, Slot 1, IOC 1"
end
```

Specifying an Adapter

An adapter is a collection of sessions. This collection is presented to the Linux kernel as if the collection was a single Fibre Channel adapter. The host system has no information regarding session connectivity. It only sees the end target fibre channel devices. The adapter section of the **`qlgc_srp`** configuration file contains multiple parameters. These parameters are listed in the adapter section of the `ib_qlgc_srp_build_cfg` script system output shown in "Determining the values to use for the configuration" on page B-6 The following example specifies an adapter:

```
adapter
begin
    description: "Oracle RAID Array"
end
```

Restarting the SRP Module

For changes to take effect, including changes to the SRP map on the VIO card, SRP will need to be restarted. To restart the `qlgc_srp` driver, use the following command:

```
/etc/init.d/qlgc_srp restart
```

Configuring an Adapter with Multiple Sessions

Each adapter can have an unlimited number of sessions attached to it. Unless `round robin` is specified, SRP will only use one session at a time. However, there is still an advantage to configuring an adapter with multiple sessions. For example, if an adapter is configured with only one session and that session fails, all SCSI I/Os on that session will fail and access to SCSI target devices will be lost. While the `qlgc_srp` module will attempt to recover the broken session, this may take some time (e.g., if a cable was pulled, the FC port has failed, or an adapter has failed). However, if the host is using an adapter configured with multiple sessions and the current session fails, the host will automatically switch to an alternate session. The result is that the host can quickly recover and continue to access the SCSI target devices.

⚠WARNING

When using two VIO hardware cards within one Adapter, the cards must have *identical* Fibre Channel configurations and maps. Data corruption can result from using different configurations and/or maps.

When the `qlgc_srp` module encounters an adapter command, that adapter is assigned all previously defined sessions (that have not been assigned to other adapters). This makes it easy to configure a system for multiple SRP adapters. The following is an example configuration that uses multiple sessions and adapters:

```
session
begin
  card: 0
  port: 2
  targetIOCProfileIdString: "FVIC in Chassis
0x00066A005000011D, Slot 1, IOC 1"
  initiatorExtension: 3
end

adapter
begin
description: "Test Device"
end

session
begin
  card: 0
  port: 1
  targetIOCProfileIdString: "FVIC in Chassis
0x00066A005000011D, Slot 2, IOC 1"
  initiatorExtension: 2
end

adapter
begin
description: "Test Device 1"
end

session
begin
  card: 0
  port: 1
  targetIOCProfileIdString: "FVIC in Chassis
0x00066A005000011D, Slot 1, IOC 2"
  initiatorExtension: 2
```

```
end

adapter
begin
description: "Test Device 1"
end
```

Configuring Fibre Channel Failover

Fibre Channel failover is essentially failing over from one session in an adapter to another session in the same adapter.

Following is a list of the different type of failover scenarios:

- Failing over from one SRP initiator port to another.
- Failing over from a port on the VIO hardware card to another port on the VIO hardware card.
- Failing over from a port on a VIO hardware card to a port on a different VIO hardware card within the same virtual I/O chassis.
- Failing over from a port on a VIO hardware card to a port on a different VIO hardware card in a different virtual I/O chassis.

Failover Configuration File 1: Failing over from one SRP Initiator port to another

In this failover configuration file, the first session (using adapter Port 1) is used to reach the SRP Target Port. If a problem is detected in this session (e.g., the IB cable on port 1 of the adapter is pulled) then the 2nd session (using adapter Port 2) will be used.

```
#  service 0: name SRP.T10:0000000000000001 id
0x0000494353535250
session
begin
  card: 0
  port: 1
  #portGuid: 0x0002c903000010f1
  initiatorExtension: 1
  targetIOCGuid: 0x00066a01e0000149
  targetIOCProfileIdString: "BC2FC in Chassis
0x0000000000000000, Slot 6, Ioc 1"
  targetPortGid: 0xfe8000000000000000066a01e0000149
  targetExtension: 0x0000000000000001
  SID: 0x0000494353535250
  IOClass: 0xff00
end

session
begin
  card: 0
  port: 2
  #portGuid: 0x0002c903000010f2
  initiatorExtension: 1
  targetIOCGuid: 0x00066a01e0000149
  targetIOCProfileIdString: "BC2FC in Chassis
0x0000000000000000, Slot 6, Ioc 1"
  targetPortGid: 0xfe8000000000000000066a01e0000149
  targetExtension: 0x0000000000000001
  SID: 0x0000494353535250
  IOClass: 0xff00
end

adapter
begin
```

```
  adapterIODepth: 1000
  lunIODepth: 16
  adapterMaxIO: 128
  adapterMaxLUNs: 512
  adapterNoConnectTimeout: 60
  adapterDeviceRequestTimeout: 2
  # set to 1 if you want round robin load balancing
  roundrobinmode: 0
  # set to 1 if you do not want target connectivity
verification
  noverify: 0
  description: "SRP Virtual HBA 0"
end
```

Failover Configuration File 2: Failing over from a port on the VIO hardware card to another port on the VIO hardware card

```
session
begin
card: 0 (InfiniServ HCA card number)
port: 1 (InfiniServ HCA port number)
targetIOCProfileIdString: "FVIC in Chassis <FRUChassisGUID>,
<Slot>, <IOC1>"
initiatorExtension: 1
end

session
begin
card: 0 (InfiniServ HCA card number)
port: 1 (InfiniServ HCA port number)
targetIOCProfileIdString: "FVIC in Chassis <FRUChassisGUID>,
<Slot>, <IOC2>"
initiatorExtension: 1 (Here the extension should be different
if using the same IOC in this adapter for FVIC, so that
separate maps can be created for each session).
 end

adapter
begin
description: “FC port Failover”
end
```

On the VIO hardware side, the following needs to be ensured:

- The target device is discovered and configured for each of the ports that is involved in the failover.
- The SRP Initiator is discovered and configured once for each different initiatorExtension.
- Each map should use a different Configured Device, e.g Configured Device 1 has the Target being discovered over FC Port 1, and Configured Device 2 has the Target being discovered over FC Port 2.)

Failover Configuration File 3: Failing over from a port on a VIO hardware card to a port on a different VIO hardware card within the same Virtual I/O chassis

```
session
begin
card: 0 (InfiniServ HCA card number)
port: 1 (InfiniServ HCA port number)
targetIOCProfileIdString: "FVIC in Chassis <FRUChassisGUID>,
<Slot1>, <IOC>"
initiatorExtension: 1
end

session
begin
card: 0 (InfiniServ HCA card number)
port: 1 (InfiniServ HCA port number)
targetIOCProfileIdString: "FVIC in Chassis <FRUChassisGUID>,
<Slot2>, <IOC>" (Slot number differs to indicate a different
VIO card)
initiatorExtension: 1 (Here the initiator extension can be the
same as in the previous definition, because the SRP map is
being defined on a different FC gateway card)
end

adapter
begin
description: “FC Port Failover”
end
```

On the VIO hardware side, the following need to be ensured on each FVIC involved in the failover:

- The target device is discovered and configured through the appropriate FC port
- The SRP Initiator is discovered and configured once for the proper initiatorExtension.
- The SRP map created for the initiator connects to the same target

Failover Configuration File 4: Failing over from a port on a VIO hardware card to a port on a different VIO hardware card in a different Virtual I/O chassis

```
session
begin
card: 0 (InfiniServ HCA card number)
port: 1 (InfiniServ HCA port number)
targetIOCProfileIdString: "FVIC in Chassis <FRUChassisGUID1>,
<Slot1>, <IOC>"
initiatorExtension: 1
end

session
begin
card: 0 (InfiniServ HCA card number)
port: 1 (InfiniServ HCA port number)
targetIOCProfileIdString: "FVIC in Chassis <FRUChassisGUID2>,
<Slot1>, <IOC>" (Chassis GUID differs to indicate a card in a
different chassis)
initiatorExtension:1 (Here the initiator extension can be the
same as in the previous definition, because the SRP map is
being defined on a different FC gateway card)
end

adapter
begin
description: “FC Port Failover”
end
```

On the VIO hardware side, the following need to be ensured on each FVIC involved in the failover:

- The target device is discovered and configured through the appropriate FC port
- The SRP Initiator is discovered and configured once for the proper initiatorExtension.
- The SRP map created for the initiator connects to the same target

Configuring Fibre Channel Load Balancing

The following examples display typical scenarios for how to configure Fibre Channel load balancing.

In the first example, traffic going to any Fibre Channel Target Device where both ports of the VIO hardware card have a valid map, are split between the two ports of the VIO hardware card. If one of the VIO hardware ports goes down, then all of the traffic will go over the remaining port that is up.

1 Adapter Port and 2 Ports on a Single VIO

```
session
 begin
        card: 0
        port: 1
        targetIOCProfileIdString: "FVIC in Chassis
0x00066A0050000123, Slot 1, IOC 1"
        initiatorExtension: 3
 end

session
 begin
        card: 0
        port: 1
        targetIOCProfileIdString: "FVIC in Chassis
0x00066A0050000123, Slot 1, IOC 2"
        initiatorExtension: 3
 end

 adapter
 begin
        description: "Test Device"
        roundrobinmode: 1
 end
```

2 Adapter Ports and 2 Ports on a Single VIO Module

In this example, traffic is load balanced between adapter Port 2/VIO hardware Port 1 and adapter Port1/VIO hardware Port 1. If one of the sessions goes down (due to an IB cable failure or an FC cable failure), all traffic will begin using the other session.

```
session
 begin
        card: 0
        port: 2
        targetIOCProfileIdString: "FVIC in Chassis
0x00066A0050000123, Slot 1, IOC 1"
        initiatorExtension: 3
 end

session
 begin
        card: 0
        port: 1
        targetIOCProfileIdString: "FVIC in Chassis
0x00066A0050000123, Slot 1, IOC 2"
        initiatorExtension: 3
 end

 adapter
 begin
        description: "Test Device"
        roundrobinmode: 1
 end
```

Using the `roundrobinmode` Parameter

In this example, the two sessions use different VIO hardware cards as well as different adapter ports. Traffic will be load-balanced between the two sessions. If there is a failure in one of the sessions (e.g., one of the VIO hardware cards is rebooted) traffic will begin using the other session.

```
session
begin
  card: 0
  port: 2
  targetIOCProfileIdString: "FVIC in Chassis
0x00066A005000011D, Slot 1, IOC 1"
  initiatorExtension: 2
end

session
begin
  card: 0
  port: 1
  targetIOCProfileIdString: "FVIC in Chassis
0x00066A005000011D, Slot 2, IOC 1"
  initiatorExtension: 2
end

adapter
begin
  description: "Test Device"
  roundrobinmode: 1
end
```

Configuring SRP for Native IB Storage

1. Review `ib_qlgc_srp_query`.

```
QLogic Corporation. Virtual HBA (SRP) SCSI Query
Application, version 1.3.0.0.1

1 IB Host Channel Adapter present in system.
HCA Card 1      : 0x0002c9020026041c
Port 1 GUID     : 0x0002c9020026041d
Port 2 GUID     : 0x0002c9020026041e

SRP Targets     :
SRP IOC Profile : Native IB Storage SRP Driver
  SRP IOC GUID  : 0x00066a01dd000021
  SRP IU SIZE   : 320
  SRP IU SG SIZE: 15
  SRP IO CLASS  : 0xff00
  service 0     : name SRP.T10:0000000000000001 id
0x0000494353535250
  service 1     : name SRP.T10:0000000000000002 id
0x0000494353535250
  service 2     : name SRP.T10:0000000000000003 id
0x0000494353535250
  service 3     : name SRP.T10:0000000000000004 id
0x0000494353535250
  Target Path(s):
   HCA 0 Port 1 0x0002c9020026041d -> Target Port GID
0xfe80000000000000066a11dd000021
   HCA 0 Port 2 0x0002c9020026041e -> Target Port GID
0xfe80000000000000066a11dd000021
```

2. Edit `/etc/sysconfig/qlgc_srp.cfg` to add this information.

```
#  service : name SRP.T10:0000000000000001 id
0x0000494353535250
session
begin
  card: 0
  port: 1
  #portGuid: 0x0002c903000010f1
  initiatorExtension: 1
  targetIOCGuid: 0x00066a01e0000149
  targetIOCProfileIdString: "Native IB Storage SRP Driver"
  targetPortGid: 0xfe80000000000000066a01e0000149
  targetExtension: 0x0000000000000001
  SID: 0x0000494353535250
  IOClass: 0x0100
end

session
begin
  card: 0
  port: 2
  #portGuid: 0x0002c903000010f2
  initiatorExtension: 1
  targetIOCGuid: 0x00066a01e0000149
  targetIOCProfileIdString: Native IB Storage SRP Driver
  targetPortGid: 0xfe80000000000000066a01e0000149
  targetExtension: 0x0000000000000001
  SID: 0x0000494353535250
  IOClass: 0x0100
end

adapter
begin
  adapterIODepth: 1000
  lunIODepth: 16
  adapterMaxIO: 128
  adapterMaxLUNs: 512
  adapterNoConnectTimeout: 60
  adapterDeviceRequestTimeout: 2
  # set to 1 if you want round robin load balancing
```

```
  roundrobinmode: 0
  # set to 1 if you do not want target connectivity
verification
  noverify: 0
  description: "SRP Virtual HBA 0"
end
Note the correlation between the output of
ib_qlgc_srp_query and qlgc_srp.cfg
Target Path(s):
   HCA 0 Port 1 0x0002c9020026041d -> Target Port GID
0xfe80000000000000066a11dd000021
   HCA 0 Port 2 0x0002c9020026041e -> Target Port GID
0xfe80000000000000066a11dd000021

qlgc_srp.cfg:

session
  begin
   . . . .
   targetIOCGuid: 0x0002C90200400098
   targetExtension: 0x0002C90200400098
  end
adapter
begin
   description: "Native IB storage"
end
```

Notes

- There is a sample configuration in `qlgc_srp.cfg`.
  - The correct `TargetExtension` must be added to session.
  - It is important to use the IOC ID method since most Profile ID strings are not guaranteed to be unique.
- Other possible parameters:
  - `initiatorExtension` may be used by the storage device to identify the host.

Additional Details

- All LUNs found are reported to the Linux SCSI mid-layer.
- Linux may need the `max_scsi_luns` (2.4 kernels) or `max_luns` (2.6 kernels) parameter configured in `scsi_mod`.

Troubleshooting

For troubleshooting information, refer to “Troubleshooting SRP Issues” on page E-9.

OFED SRP Configuration

To use OFED SRP, follow these steps:

1. Add the line `SRP_LOAD=yes` to the module list in `/etc/infiniband/openib.conf` to have it automatically loaded.
2. Discover the SRP devices on your fabric by running this command (as a root user):

   ```
   ibsrpdm
   ```

   In the output, look for lines similar to these:

   ```
   GUID:      0002c90200402c04
   ID:       LSI Storage Systems SRP Driver 200400a0b8114527
           service entries: 1
   service[  0]: 200400a0b8114527 / SRP.T10:200400A0B8114527

   GUID:      0002c90200402c0c
   ID:       LSI Storage Systems SRP Driver 200500a0b8114527
           service entries: 1
   service[  0]: 200500a0b8114527 / SRP.T10:200500A0B8114527

   GUID:      21000001ff040bf6
   ID:        Data Direct Networks SRP Target System
           service entries: 1
   service[  0]: f60b04ff01000021 / SRP.T10:21000001ff040bf6
   ```

   Note that not all the output is shown here; key elements are expected to show the match in Step 3.

3. Choose the device you want to use, and run the command again with the `-c` option (as a root user):

```
# ibsrpdm -c
id_ext=200400A0B8114527,ioc_guid=0002c90200402c04,dgid=fe
80000000000000002c90200402c05,pkey=ffff,service_id=20040
0a0b8114527
id_ext=200500A0B8114527,ioc_guid=0002c90200402c0c,dgid=fe
80000000000000002c90200402c0d,pkey=ffff,service_id=20050
0a0b8114527
id_ext=21000001ff040bf6,ioc_guid=21000001ff040bf6,dgid=fe
80000000000000021000001ff040bf6,pkey=ffff,service_id=f60b0
4ff01000021
```

4. Find the result that corresponds to the target you want, and `echo` it into the `add_target` file:

```
echo
"id_ext=21000001ff040bf6,ioc_guid=21000001ff040bf6,dgid=f
e80000000000000021000001ff040bf6,pkey=ffff,service_id=f60b
04ff01000021,initiator_ext=0000000000000001" >
/sys/class/infiniband_srp/srp-ipath0-1/add_target
```

5. Look for the newly created devices in the `/proc/partitions` file. The file will look similar to this example (the partition names may vary):

```
# cat /proc/partitions
major minor  #blocks  name
8     64  142325760 sde
8     65  142319834 sde1
8     80   71162880 sdf
8     81   71159917 sdf1
8     96      20480 sdg
8     97      20479 sdg1
```

6. Create a mount point (as root) where you will mount the SRP device. For example:

```
mkdir /mnt/targetname
mount /dev/sde1 /mnt/targetname
```

NOTE

Use `sde1` rather than `sde`. See the `mount(8)` man page for more information on creating mount points.

Notes

C Integration with a Batch Queuing System

Most cluster systems use some kind of batch queuing system as an orderly way to provide users with access to the resources they need to meet their job's performance requirements. One task of the cluster administrator is to allow users to submit MPI jobs through these batch queuing systems.

For Open MPI, there are resources at openmpi.org that document how to use the MPI with three batch queuing systems. The links to the Frequently Asked Questions (FAQs) for each of the three batch queuing system are as follows:

- Torque / PBS Pro: http://www.open-mpi.org/faq/?category=tm
- SLURM: http://www.open-mpi.org/faq/?category=slurm
- Bproc: http://www.open-mpi.org/faq/?category=bproc

In this Appendix there are two sections which deal with process and file clean-up after batch MPI/PSM jobs have completed: Clean Termination of MPI Processes and Clean-up PSM Shared Memory Files.

Clean Termination of MPI Processes

The InfiniPath software normally ensures clean termination of all MPI programs when a job ends, but in some rare circumstances an MPI process may remain alive, and potentially interfere with future MPI jobs. To avoid this problem, run a script before and after each batch job that kills all unwanted processes. QLogic does not provide such a script, but it is useful to know how to find out which processes on a node are using the QLogic interconnect. The easiest way to do this is with the `fuser` command, which is normally installed in `/sbin`.

Run these commands as a root user to ensure that all processes are reported.

```
# /sbin/fuser -v /dev/ipath
/dev/ipath:     22648m 22651m
```

In this example, processes 22648 and 22651 are using the QLogic interconnect. It is also possible to use this command (as a root user):

```
# lsof /dev/ipath
```

This command displays a list of processes using InfiniPath. Additionally, to get all processes, including `stats` programs, `ipath_sma`, diags, and others, run the program in this way:

```
# /sbin/fuser -v /dev/ipath*
```

`lsof` can also take the same form:

```
# lsof /dev/ipath*
```

The following command terminates all processes using the QLogic interconnect:

```
# /sbin/fuser -k /dev/ipath
```

For more information, see the man pages for `fuser(1)` and `lsof(8)`.

NOTE

Hard and explicit program termination, such as `kill -9` on the mpirun Process ID (PID), may result in Open MPI being unable to guarantee that the `/dev/shm` shared memory file is properly removed. As many stale files accumulate on each node, an error message can appear at startup:

```
node023:6.Error creating shared memory object in
shm_open(/dev/shm may have stale shm files that need
to be removed):
```

If this occurs, refer to Clean-up PSM Shared Memory Files for information.

Clean-up PSM Shared Memory Files

In some cases if a PSM job terminates abnormally, such as with a segmentation fault, there could be POSIX shared memory files leftover in the /dev/shm directory. The file is owned by the user and in permission `-rwx------`, it can be removed either by the user or by root.

PSM relies on the MPI implementation to cleanup after abnormal job termination. In cases where this does not occur there may be leftover share memory files. To clean up the system, create, save, and run the following PSM SHM cleanup script as root on each node. Either logon to the node, or run remote using pdsh/ssh.

```
#!/bin/sh
files=`/bin/ls /dev/shm/psm_shm.* 2> /dev/null`;
for file in $files;
do
/sbin/fuser $file > /dev/null 2>&1;
if [ $? -ne 0 ];
then
/bin/rm $file > /dev/null 2>&1;
fi;
done;
```

When the system is idle, the administrators can remove all of the shared memory files, including stale files, by using the following command:

```
# rm -rf /dev/shm/psm_shm.*
```

D Troubleshooting

This appendix describes some of the tools you can use to diagnose and fix problems. The following topics are discussed:

- Using LEDs to Check the State of the Adapter
- BIOS Settings
- Kernel and Initialization Issues
- OpenFabrics and InfiniPath Issues
- System Administration Troubleshooting
- Performance Issues
- Open MPI Troubleshooting

Troubleshooting information for hardware installation is found in the *QLogic InfiniBand*® Adapter Hardware Installation Guide and software installation is found in the *QLogic InfiniBand*® Fabric Software Installation Guide.

Using LEDs to Check the State of the Adapter

The LEDs function as link and data indicators once the InfiniPath software has been installed, the driver has been loaded, and the fabric is being actively managed by a subnet manager.

Table D-1 describes the LED states. The green LED indicates the physical link signal; the amber LED indicates the link. The green LED normally illuminates first. The normal state is *Green On*, *Amber On*. The QLE7240 and QLE7280 have an additional state, as shown in Table D-1.

Table D-1. LED Link and Data Indicators

LED States	Indication
Green OFF Amber OFF	The switch is not powered up. The software is neither installed nor started. Loss of signal. Verify that the software is installed and configured with `ipath_control -i`. If correct, check both cable connectors.

Table D-1. LED Link and Data Indicators (Continued)

LED States	Indication
Green ON Amber OFF	Signal detected and the physical link is up. Ready to talk to SM to bring the link fully up. If this state persists, the SM may be missing or the link may not be configured. Use `ipath_control -i` to verify the software state. If all IB adapters are in this state, then the SM is not running. Check the SM configuration, or install and run `opensmd`.
Green ON Amber ON	The link is configured, properly connected, and ready. Signal detected. Ready to talk to an SM to bring the link fully up. The link is configured. Properly connected and ready to receive data and link packets.
Green BLINKING (quickly) Amber ON	Indicates traffic
Green BLINKING[a] Amber BLINKING	Locates the adapter This feature is controlled by `ipath_control -b [On | Off]`

Table Notes

[a] This feature is available only on the QLE7340, QLE7342, QLE7240 and QLE7280 adapters

BIOS Settings

This section covers issues related to BIOS settings.The most important setting is Advanced Configuration and Power Interface (ACPI). This setting must be enabled. If ACPI has been disabled, it may result in initialization problems, as described in “InfiniPath Interrupts Not Working” on page D-3.

You can check and adjust the BIOS settings using the BIOS Setup utility. Check the hardware documentation that came with your system for more information.

Kernel and Initialization Issues

Issues that may prevent the system from coming up properly are described in the following sections.

Driver Load Fails Due to Unsupported Kernel

If you try to load the InfiniPath driver on a kernel that InfiniPath software does not support, the load fails. Error messages similar to this display:

```
modprobe: error inserting
'/lib/modules/2.6.3-1.1659-smp/updates/kernel/drivers/infiniband/hw/qib/ib_qib.ko': -1 Invalid module format
```

To correct this problem, install one of the appropriate supported Linux kernel versions, then reload the driver.

Rebuild or Reinstall Drivers if Different Kernel Installed

If you upgrade the kernel, then you must reboot and then rebuild or reinstall the InfiniPath kernel modules (drivers). QLogic recommends using the IFS Software Installation TUI to preform this rebuild or reinstall. Refer to the *QLogic Fabric Software Installation Guide* for more information.

InfiniPath Interrupts Not Working

The InfiniPath driver cannot configure the InfiniPath link to a usable state unless interrupts are working. Check for this problem with the command:

```
$ grep ib_qib /proc/interrupts
```

Normal output is similar to this:

```
            CPU0       CPU1
185:      364263          0   IO-APIC-level  ib_qib
```

NOTE

The output you see may vary depending on board type, distribution, or update level.

If there is no output at all, the driver initialization failed. For more information on driver problems, see “Driver Load Fails Due to Unsupported Kernel” on page D-3 or “InfiniPath ib_qib Initialization Failure” on page D-5.

If the output is similar to one of these lines, then interrupts are not being delivered to the driver.

```
66: 0 0 PCI-MSI  ib_qib
185:00IO-APIC-level  ib_qib
```

The following message appears when driver has initialized successfully, but no interrupts are seen within 5 seconds.

```
ib_qib 0000:82:00.0: No interrupts detected.
```

A zero count in all CPU columns means that no InfiniPath interrupts have been delivered to the processor.

The possible causes of this problem are:

- Booting the Linux kernel with ACPI disabled on either the boot command line or in the BIOS configuration
- Other `infinipath` initialization failures

To check if the kernel was booted with the `noacpi` or `pci=noacpi` option, use this command:

```
$ grep -i acpi /proc/cmdline
```

If output is displayed, fix the kernel boot command line so that ACPI is enabled. This command line can be set in various ways, depending on your distribution. If no output is displayed, check that ACPI is enabled in your BIOS settings.

To track down other initialization failures, see “InfiniPath ib_qib Initialization Failure” on page D-5.

The program `ipath_checkout` can also help flag these kinds of problems. See “ipath_checkout” on page G-25 for more information.

OpenFabrics Load Errors if `ib_qib` Driver Load Fails

When the `ib_qib` driver fails to load, the other OpenFabrics drivers/modules will load and be shown by `lsmod`, but commands like `ibstatus`, `ibv_devinfo`, and `ipath_control -i` will fail as follows:

```
# ibstatus
Fatal error:  device '*': sys files not found
(/sys/class/infiniband/*/ports)
# ibv_devinfo
libibverbs: Fatal: couldn't read uverbs ABI version.
No IB devices found
# ipath_control -i
InfiniPath driver not loaded ?
No InfiniPath info available
```

InfiniPath ib_qib Initialization Failure

There may be cases where ib_qib was not properly initialized. Symptoms of this may show up in error messages from an MPI job or another program. Here is a sample command and error message:

```
$ mpirun -np 2 -m ~/tmp/mbu13 osu_latency
<nodename>:ipath_userinit: assign_port command failed:
Network is down
<nodename>:can't open /dev/ipath, network down
```

This will be followed by messages of this type after 60 seconds:

```
MPIRUN<node_where_started>: 1 rank has not yet exited 60
seconds after rank 0 (node <nodename>) exited without reaching
MPI_Finalize().
MPIRUN<node_where_started>:Waiting at most another 60 seconds
for the remaining ranks to do a clean shutdown before
terminating 1 node processes.
```

If this error appears, check to see if the InfiniPath driver is loaded by typing:

```
$ lsmod | grep ib_qib
```

If no output is displayed, the driver did not load for some reason. In this case, try the following commands (as root):

```
# modprobe -v ib_qib
# lsmod | grep ib_qib
# dmesg | grep -i ib_qib | tail -25
```

The output will indicate whether the driver has loaded. Printing out messages using dmesg may help to locate any problems with ib_qib.

If the driver loaded, but MPI or other programs are not working, check to see if problems were detected during the driver and QLogic hardware initialization with the command:

```
$ dmesg | grep -i ib_qib
```

This command may generate more than one screen of output.

Also, check the link status with the commands:

```
$ cat /sys/class/infiniband/ipath*/device/status_str
```

These commands are normally executed by the ipathbug-helper script, but running them separately may help locate the problem.

See also "status_str" on page G-35 and "ipath_checkout" on page G-25.

MPI Job Failures Due to Initialization Problems

If one or more nodes do not have the interconnect in a usable state, messages similar to the following appear when the MPI program is started:

```
userinit: userinit ioctl failed: Network is down [1]: device
init failed
userinit: userinit ioctl failed: Fatal Error in
keypriv.c(520): device init failed
```

These messages may indicate that a cable is not connected, the switch is down, SM is not running, or that a hardware error occurred.

OpenFabrics and InfiniPath Issues

The following sections cover issues related to OpenFabrics (including Subnet Managers) and InfiniPath.

Stop Infinipath Services Before Stopping/Restarting InfiniPath

The following Infinipath services must be stopped before stopping/starting/restarting InfiniPath:

- QLogic Fabric Manager
- OpenSM
- SRP

Here is a sample command and the corresponding error messages:

```
# /etc/init.d/openibd stop
Unloading infiniband modules: sdp cm umad uverbs ipoib sa
ipath mad coreFATAL:Module ib_umad is in use.
Unloading infinipath modules FATAL: Module ib_qib is in use.
[FAILED]
```

Manual Shutdown or Restart May Hang if NFS in Use

If you are using NFS over IPoIB and use the manual `/etc/init.d/openibd stop` (or `restart`) command, the shutdown process may silently hang on the `fuser` command contained within the script. This is because `fuser` cannot traverse down the tree from the mount point once the mount point has disappeared. To remedy this problem, the `fuser` process itself needs to be killed. Run the following command either as a root user or as the user who is running the `fuser` process:

```
# kill -9 fuser
```

The shutdown will continue.

This problem is not seen if the system is rebooted or if the filesystem has already been unmounted before stopping `infinipath`.

Load and Configure IPoIB Before Loading SDP

SDP generates Connection Refused errors if it is loaded before IPoIB has been loaded and configured. To solve the problem, load and configure IPoIB first.

Set `$IBPATH` for OpenFabrics Scripts

The environment variable `$IBPATH` must be set to `/usr/bin`. If this has not been set, or if you have it set to a location other than the installed location, you may see error messages similar to the following when running some OpenFabrics scripts:

```
/usr/bin/ibhosts: line 30: /usr/local/bin/ibnetdiscover: No
such file or directory
```

For the OpenFabrics commands supplied with this InfiniPath release, set the variable (if it has not been set already) to `/usr/bin`, as follows:

```
$ export IBPATH=/usr/bin
```

SDP Module Not Loading

If the settings for debug level and the zero copy threshold from InfiniPath release 2.0 are present in the release 2.2 `/etc/modprobe.conf` file (RHEL) or `/etc/modprobe.conf.local` (SLES) file, the SDP module may not load:

```
options ib_sdp sdp_debug_level=4
sdp_zcopy_thrsh_src_default=10000000
```

To solve the problem, remove this line.

`ibsrpdm` Command Hangs when Two Host Channel Adapters are Installed but Only Unit 1 is Connected to the Switch

If multiple IB adapters (unit 0 and unit 1) are installed and only unit 1 is connected to the switch, the `ibsrpdm` command (to set up an SRP target) can hang. If unit 0 is connected and unit 1 is disconnected, the problem does not occur.

When only unit 1 is connected to the switch, use the `-d` option with `ibsrpdm`. Then, using the output from the `ibsrpdm` command, echo the new target information into `/sys/class/infiniband_srp/srp-ipath1-1/add_target`.

For example:

```
# ibsrpdm -d /dev/infiniband/umad1 -c
# echo \
id_ext=21000001ff040bf6,ioc_guid=21000001ff040bf6,dgid=fe8000
0000000000021000001ff040bf6,pkey=ffff,service_id=f60b04ff01000
021 > /sys/class/infiniband_srp/srp-ipath0-1/add_target
```

Outdated ipath_ether Configuration Setup Generates Error

Ethernet emulation (`ipath_ether`) has been removed in this release, and, as a result, an error may be seen if the user still has an alias set previously by `modprobe.conf` (for example, `alias eth2 ipath_ether`).

When `ifconfig` or `ifup` are run, the error will look similar to this (assuming `ipath_ether` was used for `eth2`):

```
eth2: error fetching interface information: Device not found
```

To prevent the error message, remove the following files (assuming `ipath_ether` was used for `eth2`):

`/etc/sysconfig/network-scripts/ifcfg-eth2` (for RHEL)

`/etc/sysconfig/network/ifcfg-eth2` (for SLES)

QLogic recommends using the IP over IB protocol (IPoIB-CM), included in the standard OpenFabrics software releases, as a replacement for `ipath_ether`.

System Administration Troubleshooting

The following sections provide details on locating problems related to system administration.

Broken Intermediate Link

Sometimes message traffic passes through the fabric while other traffic appears to be blocked. In this case, MPI jobs fail to run.

In large cluster configurations, switches may be attached to other switches to supply the necessary inter-node connectivity. Problems with these inter-switch (or intermediate) links are sometimes more difficult to diagnose than failure of the final link between a switch and a node. The failure of an intermediate link may allow some traffic to pass through the fabric while other traffic is blocked or degraded.

If you notice this behavior in a multi-layer fabric, check that all switch cable connections are correct. Statistics for managed switches are available on a per-port basis, and may help with debugging. See your switch vendor for more information.

QLogic recommends using FastFabric to help diagnose this problem. If FastFabric is not installed in the fabric, there are two diagnostic tools, `ibhosts` and `ibtracert`, that may also be helpful. The tool `ibhosts` lists all the IB nodes that the subnet manager recognizes. To check the IB path between two nodes, use the `ibtracert` command.

Performance Issues

The following sections discuss known performance issues.

Large Message Receive Side Bandwidth Varies with Socket Affinity on Opteron Systems

On Opteron systems, when using the QLE7240 or QLE7280 in DDR mode, there is a receive side bandwidth bottleneck for CPUs that are not adjacent to the PCI Express root complex. This may cause performance to vary. The bottleneck is most obvious when using SendDMA with large messages on the farthest sockets. The best case for SendDMA is when both sender and receiver are on the closest sockets. Overall performance for PIO (and smaller messages) is better than with SendDMA.

Erratic Performance

Sometimes erratic performance is seen on applications that use interrupts. An example is inconsistent SDP latency when running a program such as `netperf`. This may be seen on AMD-based systems using the QLE7240 or QLE7280 adapters. If this happens, check to see if the program `irqbalance` is running. This program is a Linux daemon that distributes interrupts across processors. However, it may interfere with prior interrupt request (IRQ) affinity settings, introducing timing anomalies. After stopping this process (as a root user), bind IRQ to a CPU for more consistent performance. First, stop `irqbalance`:

```
# /sbin/chkconfig irqbalance off
# /etc/init.d/irqbalance stop
```

Next, find the IRQ number and bind it to a CPU. The IRQ number can be found in one of two ways, depending on the system used. Both methods are described in the following paragraphs.

NOTE

Take care when cutting and pasting commands from PDF documents, as quotes are special characters and may not be translated correctly.

Method 1

Check to see if the IRQ number is found in `/proc/irq/`*`xxx`*, where *`xxx`* is the IRQ number in `/sys/class/infiniband/ipath*/device/irq`. Do this as a root user. For example:

```
# my_irq=`cat /sys/class/infiniband/ipath*/device/irq`
# ls /proc/irq
```

If `$my_irq` can be found under `/proc/irq/`, then type:

```
# echo 01 > /proc/irq/$my_irq/smp_affinity
```

Method 2

If command from Method 1, `ls /proc/irq`, cannot find `$my_irq`, then use the following commands instead:

```
# my_irq=`cat /proc/interrupts|grep ib_qib|awk \
'{print $1}'|sed -e 's/://'`
# echo 01 > /proc/irq/$my_irq/smp_affinity
```

This method is not the first choice because, on some systems, there may be two rows of ib_qib output, and you will not know which one of the two numbers to choose. However, if you cannot find $my_irq listed under /proc/irq (Method 1), this type of system most likely has only one line for ib_qib listed in /proc/interrupts, so you can use Method 2.

Here is an example:

```
# cat /sys/class/infiniband/ipath*/device/irq
98
# ls /proc/irq
0  10   11  13  15  233  4  50  7  8  90
1  106  12  14  2   3    5  58  66  74 9
```

(Note that you cannot find 98.)

```
# cat /proc/interrupts|grep ib_qib|awk \
'{print $1}'|sed -e 's/://'
106
# echo 01 > /proc/irq/106/smp_affinity
```

Using the echo command immediately changes the processor affinity of an IRQ.

NOTE

- The contents of the smp_affinity file may not reflect the expected values, even though the affinity change has taken place.
- If the driver is reloaded, the affinity assignment will revert to the default, so you will need to reset it to the desired value.

You can look at the stats in /proc/interrupts while the adapter is active to observe which CPU is fielding ib_qib interrupts.

Immediately change the processor affinity of an IRQ

To immediately change the processor affinity of an IRQ, execute a command similar to the following, as a root user:

```
echo 01 > /proc/irq/$my_irq/smp_affinity
```

The contents of the smp_affinity file may not reflect the expected values, even though the affinity change has taken place. If the driver is reloaded, the affinity assignment will revert to the default, so you will need to reset it to the desired value. Look at the stats in /proc/interrupts while the adapter is active to see which CPU is fielding ib_qib interrupts.

Performance Warning if ib_qib Shares Interrupts with eth0

When ib_qib shares interrupts with eth0, performance may be affected the OFED ULPs, such as IPoIB. A warning message appears in syslog, and also on the console or tty session where /etc/init.d/openibd start is run (if messages are set up to be displayed). Messages are in this form:

```
Nov  5 14:25:43 <nodename> infinipath: Shared interrupt will
affect performance: vector 169: devices eth0, ib_qib
```

Check /proc/interrupts: "169" is in the first column, and "devices" are shown in the last column.

You can also contact your system vendor to see if the BIOS settings can be changed to avoid the problem.

Open MPI Troubleshooting

Problems specific to compiling and running Open MPI programs are described in the following sections.

Invalid Configuration Warning

Open MPI warns about a invalid configuration every time it is run with the following warning:

```
WARNING: There are more than one active ports on host
'st2107', but the default subnet GID prefix was detected on
more than one of these ports.  If these ports are connected to
different physical IB networks, this configuration will fail
in Open MPI.  This version of Open MPI requires that every
physically separate IB subnet that is used between connected
MPI processes must have different subnet ID values.
```

When connecting 2 ports of an HCA to different fabrics, it is a mandatory requirement that the SubnetPrefix for those two fabrics be different and non-default (for example, not FE80000000000000) based on the Fabric Manager configuration file. The config_generate tool for the Fabric Manager will help generate such files. Refer to the *QLogic Fabric Manager User Guide* for more information about the config_generate tool.

E ULP Troubleshooting

Troubleshooting VirtualNIC and VIO Hardware Issues

To verify that an IB host can access an Ethernet system through the EVIC, issue a **`ping`** command to the Ethernet system from the IB host. Make certain that the route to the Ethernet system is using the VIO hardware by using the Linux **`route`** command on the IB host, then verify that the route to the subnet is using one of the virtual Ethernet interfaces (i.e., an EIOC).

NOTE

If the **`ping`** command fails, check the following:

- The logical connection between the IB host and the EVIC Checking the logical connection between the IB Host and the VIO hardware.
- The interface definitions on the host Checking the interface definitions on the host.
- The physical connection between the VIO hardware and the Ethernet network Verify the physical connection between the VIO hardware and the Ethernet network.

Checking the logical connection between the IB Host and the VIO hardware

To determine if the logical connection between the IB host and the VIO hardware is correct, check the following:

- The correct VirtualNIC driver is running.
- The **`/etc/infiniband/qlgc_vnic.cfg`** file contains the desired information.
- The host can communicate with the I/O Controllers (IOCs) of the VIO hardware.

Verify that the proper VirtualNIC driver is running

Check that a VirtualNIC driver is running by issuing an **lsmod** command on the IB host. Make sure that the **qlgc_vnic** is displayed on the list of modules. Following is an example:

```
st186:~ # lsmod
Module                  Size  Used by
cpufreq_ondemand       25232  1
cpufreq_userspace      23552  0
cpufreq_powersave      18432  0
powernow_k8            30720  2
freq_table             22400  1 powernow_k8
qlgc_srp               93876  0
qlgc_vnic             116300  0#
```

Verifying that the qlgc_vnic.cfg file contains the correct information

Use the following scenarios to verify that the **qlgc_vnic.cfg** file contains a definition for the applicable virtual interface:

Issue the command **ib_qlgc_vnic_query** to get the list of IOCs the host can see.

If the list is empty, there may be a syntax error in the **qlgc_vnic.cfg** file (e.g., a missing semicolon). Look in **/var/log/messages** at the time **qlgc_vnic** was last started to see if any error messages were put in the log at that time.

If the **qlgc_vnic.cfg** file has been edited since the last time the VirtualNIC driver was started, the driver needs restarted. To restart the driver, so that it uses the current **qlgc_vnic.cfg** file, issue a **/etc/init.d/qlgc_vnic restart**.

Verifying that the host can communicate with the I/O Controllers (IOCs) of the VIO hardware

To display the Ethernet VIO cards that the host can see and communicate with, issue the command **`ib_qlgc_vnic_query`**. The system returns information similar to the following:

```
IO Unit Info:
    port LID:        0003
    port GID:        fe80000000000000066a0258000001
    change ID:       0009
    max controllers: 0x03

    controller[  1]
        GUID:      00066a0130000001
        vendor ID: 00066a
        device ID: 000030
        IO class : 2000
        ID:        Chassis 0x00066A00010003F2, Slot 1, IOC 1
        service entries: 2
            service[  0]: 1000066a00000001 /
InfiniNIC.InfiniConSys.Control:01
            service[  1]: 1000066a00000101 /
InfiniNIC.InfiniConSys.Data:01

    controller[  2]
        GUID:      00066a0230000001
        vendor ID: 00066a
        device ID: 000030
        IO class : 2000
        ID:        Chassis 0x00066A00010003F2, Slot 1, IOC 2
        service entries: 2
            service[  0]: 1000066a00000002 /
InfiniNIC.InfiniConSys.Control:02
            service[  1]: 1000066a00000102 /
InfiniNIC.InfiniConSys.Data:02

    controller[  3]
        GUID:      00066a0330000001
        vendor ID: 00066a
        device ID: 000030
        IO class : 2000
```

```
        ID:        Chassis 0x00066A00010003F2, Slot 1, IOC 3
        service entries: 2
            service[  0]: 1000066a00000003 /
InfiniNIC.InfiniConSys.Control:03
            service[  1]: 1000066a00000103 /
InfiniNIC.InfiniConSys.Data:03
```

When **ib_qlgc_vnic_query** is run with **-e** option, it reports the IOCGUID information. With the **-s** option it reports the IOCSTRING information for the Virtual I/O hardware IOCs present on the fabric. Following is an example:

```
# ib_qlgc_vnic_query -e
ioc_guid=00066a0130000001,dgid=fe80000000000000066a02580000
01,pkey=ffff
ioc_guid=00066a0230000001,dgid=fe80000000000000066a02580000
01,pkey=ffff
ioc_guid=00066a0330000001,dgid=fe80000000000000066a02580000
01,pkey=ffff
#ib_qlgc_vnic_query -s
"Chassis 0x00066A00010003F2, Slot 1, IOC 1"
"Chassis 0x00066A00010003F2, Slot 1, IOC 2"
"Chassis 0x00066A00010003F2, Slot 1, IOC 3"
#ib_qlgc_vnic_query -es
ioc_guid=00066a0130000001,dgid=fe80000000000000066a02580000
01,pkey=ffff,"Chassis 0x00066A00010003F2, Slot 1, IOC 1"
ioc_guid=00066a0230000001,dgid=fe80000000000000066a02580000
01,pkey=ffff,"Chassis 0x00066A00010003F2, Slot 1, IOC 2"
ioc_guid=00066a0330000001,dgid=fe80000000000000066a02580000
01,pkey=ffff,"Chassis 0x00066A00010003F2, Slot 1, IOC 3"
```

If the host can not see applicable IOCs, there are two things to check. First, verify that the adapter port specified in the **eioc** definition of the **/etc/infiniband/qlgc_vnic.cfg** file is active. This is done using the **ibv_devinfo** commands on the host, then checking the value of **state**. If the state is not **Port_Active**, the adapter port is not logically connected to the fabric. It is possible that one of the adapter ports is not physically connected to an IB switch. For example:

```
st139:~ # ibv_devinfo
hca_id: mlx4_0
        fw_ver:                         2.2.000
        node_guid:                      0002:c903:0000:0f80
        sys_image_guid:                 0002:c903:0000:0f83
        vendor_id:                      0x02c9
        vendor_part_id:                 25418
        hw_ver:                         0xA0
        board_id:                       MT_04A0110002
        phys_port_cnt:                  2
                port:   1
                        state:                  PORT_ACTIVE (4)
                        max_mtu:                2048 (4)
                        active_mtu:             2048 (4)
                        sm_lid:                 1
                        port_lid:               8
                        port_lmc:               0x00

                port:   2
                        state:                  PORT_ACTIVE (4)
                        max_mtu:                2048 (4)
                        active_mtu:             2048 (4)
                        sm_lid:                 1
                        port_lid:               9
                        port_lmc:               0x00#
```

Second, verify that the adapter port specified in the EIOC definition is the correct port. The host sees the IOCs, but not over the adapter port in the definition of the IOC. For example, the host may see the IOCs over adapter Port 1, but the **eioc** definition in the **/etc/infiniband/qlgc_vnic.cfg** file specifies **PORT=2**.

Another reason why the host might not be able to see the necessary IOCs is that the subnet manager has gone down. Issue an **iba_saquery** command to make certain that the response shows all of the nodes in the fabric. If an error is returned and the adapter is physically connected to the fabric, then the subnet manager has gone down, and this situation needs to be corrected.

Checking the interface definitions on the host

If it is not possible to ping from an IB host to the Ethernet host, and the **ViPort State** of the interface is **VIPORT_CONNECTED**, then issue an **ifconfig** command. The interfaces defined in the configuration files listed in **/etc/sysconfig/network** directory for SLES hosts or the **/etc/sysconfig/network-scripts** for Red Hat hosts should be displayed in the list of interfaces in the **ifconfig** output. For example, the **ifconfig** file should show an interface for each EIOC configuration file in the following list:

```
# ls /etc/sysconfig/network-scripts
ifcfg-eioc1
ifcfg-eioc2
ifcfg-eioc3
ifcfg-eioc4
ifcfg-eioc5
ifcfg-eioc6
```

Interface does not show up in output of 'ifconfig'

If an interface is not displayed in the output of an **ifconfig** command, there is most likely a problem in the definition of that interface in the **/etc/sysconfig/network-scripts/ifcfg-<NAME>** (for RedHat systems) or **/etc/sysconfig/network/ifcfg-<NAME>** (for SuSE systems) file, where **<NAME>** is the name of the virtual interface (e.g., **eioc1**).

NOTE

For the remainder of this section, **ifcfg directory** refers to **/etc/sysconfig/network-scripts/** on RedHat systems, and **/etc/sysconfig/network** on SuSE systems.

Issue an **ifup <NAME>** command. If the interface is displayed when issuing an **ifconfig** command, there may be a problem with the way the interface startup is defined in the **ifcfg directory'/ifcfg-<NAME>** file that is preventing the interface from coming up automatically.
If the interface does not come up, check the interface definitions in the **ifcfg** directory. Make certain that there are no misspellings in the **ifcfg-<NAME>** file.

Example of ifcfg-eiocx setup for RedHat systems:

DEVICE=eioc1

BOOTPROTO=static

IPADDR=172.26.48.132

BROADCAST=172.26.63.130

NETMASK=255.255.240.0

NETWORK=172.26.48.0

ONBOOT=yes

TYPE=Ethernet

Example of ifcfg-eiocx setup for SuSE and SLES systems:

BOOTPROTO='static'

IPADDR='172.26.48.130'

BROADCAST='172.26.63.255'

NETMASK='255.255.240.0'

NETWORK='172.26.48.0'

STARTMODE='hotplug'

TYPE='Ethernet'

Verify the physical connection between the VIO hardware and the Ethernet network

If the interface is displayed in an **`ifconfig`** and a ping between the IB host and the Ethernet host is still unsuccessful, verify that the VIO hardware Ethernet ports are physically connected to the correct Ethernet network. Verify that the Ethernet port corresponding to the **`IOCGUID`** for the interface to be used is connected to the expected Ethernet network.

There are up to 6 IOC GUIDs on each VIO hardware module (6 for the IB/Ethernet Bridge Module, 2 for the EVIC), one for each Ethernet port. If a VIO hardware module can be seen from a host, the **`ib_qlgc_vnic_query -s`** file displays information similar to:

```
EVIC in Chassis 0x00066a000300012a, Slot 19, Ioc 1
EVIC in Chassis 0x00066a000300012a, Slot 19, Ioc 2
EVIC in Chassis 0x00066a000300012a, Slot 8, Ioc 1
EVIC in Chassis 0x00066a000300012a, Slot 8, Ioc 2
EVIC in Chassis 0x00066a00da000100, Slot 2, Ioc 1
EVIC in Chassis 0x00066a00da000100, Slot 2, Ioc 2
```

Troubleshooting SRP Issues

`ib_qlgc_srp_stats` showing session in disconnected state

Problem:

If the session is part of a multi-session adapter, **`ib_qlgc_srp_stats`** will show it to be in the disconnected state. For example:

```
SCSI Host #        : 17                  | Mode             :
ROUNDROBIN
       Trgt Adapter Depth : 1000             | Verify Target
: Yes
       Rqst Adapter Depth : 1000             | Rqst LUN Depth
: 16
       Tot Adapter Depth  : 1000             | Tot LUN Depth
: 16
       Act Adapter Depth  : 998              | Act LUN Depth
: 16
       Max LUN Scan       : 512              | Max IO
: 131072 (128 KB)
       Max Sectors        : 256              | Max SG Depth
: 33
        Session Count      : 2                  | No Connect
T/O  : 60 Second(s)
       Register In Order  : ON               | Dev Reqst T/O
: 2 Second(s)
        Description        : SRP Virtual HBA 1


       Session           : Session 1         | State
: Disconnected
       Source GID         : 0xfe8000000000000066a000100d051
        Destination GID    :
0xfe8000000000000066a0260000165
        SRP IOC Profile    : Chassis 0x00066A0001000481, Slot
1, IOC 1
        SRP Target IOClass : 0xFF00             | SRP Target
SID  : 0x0000494353535250
        SRP IPI Guid       : 0x00066a000100d051 | SRP IPI
Extnsn  : 0x0000000000000001
        SRP TPI Guid       : 0x00066a0138000165 | SRP TPI
Extnsn  : 0x0000000000000001
       Source LID         : 0x000b             | Dest LID
: 0x0004
      Completed Sends    : 0x00000000000002c0 | Send Errors
```

```
: 0x0000000000000000
        Completed Receives : 0x00000000000002c0 | Receive
Errors  : 0x0000000000000000
      Connect Attempts   : 0x0000000000000000 | Test Attempts
: 0x0000000000000000
        Total SWUs         : 0x00000000000003e8 | Available
SWUs  : 0x00000000000003e8
       Busy SWUs         : 0x0000000000000000 | SRP Req Limit
: 0x00000000000003e8
       SRP Max ITIU      : 0x0000000000000140 | SRP Max TIIU
: 0x0000000000000140
        Host Busys         : 0x0000000000000000 | SRP Max SG
Used : 0x000000000000000f

       Session           : Session 2          | State
: Disconnected
       Source GID         : 0xfe8000000000000000066a000100d052
        Destination GID    :
0xfe8000000000000000066a0260000165
        SRP IOC Profile    : Chassis 0x00066A0001000481, Slot
1, IOC 2
        SRP Target IOClass : 0xFF00             | SRP Target
SID  : 0x0000494353535250
        SRP IPI Guid       : 0x00066a000100d052 | SRP IPI
Extnsn  : 0x0000000000000001
        SRP TPI Guid       : 0x00066a0238000165 | SRP TPI
Extnsn  : 0x0000000000000001
       Source LID         : 0x000c             | Dest LID
: 0x0004
      Completed Sends    : 0x00000000000001c8 | Send Errors
: 0x0000000000000000
        Completed Receives : 0x00000000000001c8 | Receive
Errors  : 0x0000000000000000
      Connect Attempts   : 0x0000000000000000 | Test Attempts
: 0x0000000000000000
        Total SWUs         : 0x00000000000003e8 | Available
SWUs  : 0x00000000000003e8
       Busy SWUs         : 0x0000000000000000 | SRP Req Limit
: 0x00000000000003e8
       SRP Max ITIU      : 0x0000000000000140 | SRP Max TIIU
: 0x0000000000000140
        Host Busys         : 0x0000000000000000 | SRP Max SG
Used : 0x000000000000000f
```

Solution:

Perhaps an interswitch cable has been disconnected, or the VIO hardware is offline, or the Chassis/Slot does not contain a VIO hardware card. Instead of looking at this file, use the **`ib_qlgc_srp_query`** command to verify that the desired adapter port is in the active state.

NOTE

It is normal to see the "**`Can not find a path`**" message when the system first boots up. Sometimes SRP comes up before the subnet manager has brought the port state of the adapter port to active. If the adapter port is not active, SRP will not be able to find the VIO hardware card. Use the appropriate OFED command to show the port state.

Session in 'Connection Rejected' state

Problem:

The session is in the '**`Connection Rejected`**' state according to **`/var/log/messages`**. If the session is part of a multi-session adapter, **`ib_qlgc_srp_stats`** shows it in the "**`Connection Rejected`**" state.

A host displays

```
"Connection Failed for Session X: IBT Code = 0x0
"Connection Failed for Session X: SRP Code = 0x1003
"Connection Rejected"
```

Following is an example:

```
SCSI Host #        : 17                  | Mode            :
ROUNDROBIN
       Trgt Adapter Depth : 1000             | Verify Target
: Yes
       Rqst Adapter Depth : 1000             | Rqst LUN Depth
: 16
       Tot Adapter Depth  : 1000             | Tot LUN Depth
: 16
       Act Adapter Depth  : 998              | Act LUN Depth
: 16
       Max LUN Scan       : 512              | Max IO
: 131072 (128 KB)
       Max Sectors        : 256              | Max SG Depth
: 33
        Session Count      : 2                  | No Connect
T/O  : 60 Second(s)
       Register In Order  : ON               | Dev Reqst T/O
: 2 Second(s)
        Description        : SRP Virtual HBA 1


       Session           : Session 1         | State
: Disconnected
       Source GID         : 0xfe8000000000000066a000100d051
        Destination GID    :
0xfe8000000000000066a0260000165
        SRP IOC Profile    : Chassis 0x00066A0001000481, Slot
1, IOC 1
        SRP Target IOClass : 0xFF00             | SRP Target
SID  : 0x0000494353535250
        SRP IPI Guid       : 0x00066a000100d051 | SRP IPI
Extnsn  : 0x0000000000000001
        SRP TPI Guid       : 0x00066a0138000165 | SRP TPI
Extnsn  : 0x0000000000000001
       Source LID         : 0x000b             | Dest LID
: 0x0004
      Completed Sends    : 0x00000000000002c0 | Send Errors
: 0x0000000000000000
        Completed Receives : 0x00000000000002c0 | Receive
Errors  : 0x0000000000000000
      Connect Attempts   : 0x0000000000000000 | Test Attempts
: 0x0000000000000000
        Total SWUs         : 0x00000000000003e8 | Available
```

```
SWUs  : 0x00000000000003e8
       Busy SWUs         : 0x0000000000000000 | SRP Req Limit
: 0x00000000000003e8
       SRP Max ITIU      : 0x0000000000000140 | SRP Max TIIU
: 0x0000000000000140
        Host Busys         : 0x0000000000000000 | SRP Max SG
Used : 0x000000000000000f

       Session           : Session 2          | State
: Disconnected
       Source GID         : 0xfe80000000000000066a000100d052
        Destination GID    :
0xfe80000000000000066a0260000165
        SRP IOC Profile    : Chassis 0x00066A0001000481, Slot
1, IOC 2
        SRP Target IOClass : 0xFF00             | SRP Target
SID  : 0x0004494353535250
        SRP IPI Guid       : 0x00066a000100d052 | SRP IPI
Extnsn  : 0x0000000000000001
        SRP TPI Guid       : 0x00066a0238000165 | SRP TPI
Extnsn  : 0x0000000000000001
       Source LID         : 0x000c             | Dest LID
: 0x0004
      Completed Sends    : 0x00000000000001c8 | Send Errors
: 0x0000000000000000
        Completed Receives : 0x00000000000001c8 | Receive
Errors  : 0x0000000000000000
      Connect Attempts   : 0x0000000000000000 | Test Attempts
: 0x0000000000000000
        Total SWUs         : 0x00000000000003e8 | Available
SWUs  : 0x00000000000003e8
       Busy SWUs         : 0x0000000000000000 | SRP Req Limit
: 0x00000000000003e8
       SRP Max ITIU      : 0x0000000000000140 | SRP Max TIIU
: 0x0000000000000140
        Host Busys         : 0x0000000000000000 | SRP Max SG
Used : 0x000000000000000f
```

AND

The VIO hardware displays "**Initiator Not Configured within IOU: initiator <hca port guid> <initiator extension> <target> <VIO hardware IOC GUID> <0x00000000000000001> Initiatorport identifier (<hca port guid><Initiator extension> is invalid/not allowed to use this FCIOU**".

Solution 1:

The host initiator has not been configured as an SRP initiator on the VIO hardware SRP Initiator Discovery screen. Via Chassis Viewer, bring up the SRP Initiator Discovery screen and either

Click on '**`Add New`**' to add a wildcarded entry with the initiator extension to match what is in the session entry in the **`qlgc_srp.cfg`** file, or

Click on the **`Start`** button to discover the adapter port GUID, and then click '**`Configure`**' on the row containing the adapter port GUID and give the entry a name.

Solution 2:

Check the SRP map on the VIO hardware specified in the failing Session block of the **`qlgc_srp.cfg`** file. Make certain there is a map defined for the row specified by either the **`initiatorExtension`** in the failing Session block of the **`qlgc_srp.cfg`** file or the adapter port GUID specified in the failing Session block of the **`qlgc_srp.cfg`** file. Additionally, make certain that the map in that row is in the column of the IOC specified in the failing Session block of the **`qlgc_srp.cfg`** file.

Attempts to read or write to disk are unsuccessful

Problem:

Attempts to read or write to the disk are unsuccessful when SRP comes up. About every five seconds the VIO hardware displays

"fcIOStart Failed",

"CMDisconnect called for Port: xxxxxxxx Initiator: <hca port guid> <initiator extension> Target: <VIO hardware Port GUID> <0x0000000000000001>",

and

"Target Port Deleted for Port: xxxxxxxx Initiator: <hca port guid> <initiator extension> Target: <VIO hardware Port GUID> <0x0000000000000001>"

The host log shows a session transitioning between Connected and Down. The host log also displays "**`Test Unit Ready has FAILED", "Abort Task has SUCCEEDED", and "Clear Task has FAILED`**".

Solution:

This indicates a problem in the path between the VIO hardware and the target storage device. After an SRP host has connected to the VIO hardware successfully, the host sends a "Test **Unit Ready**" command to the storage device. After five seconds, if that command is not responded to, the SRP host brings down the session and retries in five seconds. Verify that the status of the connection between the appropriate VIO hardware port and the target device is **UP** on the FCP Device Discovery screen.

Problem:

Attempts to read or write to the disk are unsuccessful, when they were previously successful. The host displays '**Sense Data indicates recovery is necessary on Session**' and the "**Test Unit Ready has FAILED**", "**Abort Task has SUCCEEDED**", "**Clear Task has FAILED**" messages.

Solution:

If there is a problem with communication between the VIO hardware and the storage device (e.g., the cable between the storage device and the Fibre Channel switch was pulled) the VIO hardware log will display a "**Connection Lost to NPort Id**" message. The next time the host tries to do an input/output (I/O), the '**Sense Data indicates recovery is necessary**' appears. Then SRP will recycle the session. As part of trying to move the session from '**Connected**' to '**Active**', SRP will issue the '**Test Unit Ready**' command.

Verify that the status of the connection between the appropriate VIO hardware port and the target device is **UP** on the FCP Device Discovery screen.

Additionally, there may occasionally be messages in the log such as:

```
Connection Failed for Session X: IBT Code = 0x0
Connection Failed for Session X: SRP Code = 0x0
```

That may indicate a problem in the path between the VIO hardware and the target storage device.

Four sessions in a round-robin configuration are active

Problem:

Four sessions in a round-robin configuration are active according to **ib_qlgc_srp_stats**. However, only one disk can be seen, although five should be seen.

Solution 1:

Make certain that **Max LUNs Scanned** is reporting the same value as **adapterMaxLUNs** is set to in **qlgc_srp.cfg**.

Solution 2:

Make certain that all sessions have a map to the same disk defined. The fact that the session is active means that the session can see a disk. However, if one of the sessions is using a map with the 'wrong' disk, then the round-robin method could lead to a disk or disks not being seen.

Which port does a port GUID refer to?

Solution:

A QLogic HCA Port GUID is of the form **00066appa0iiiiii**

where **pp** gives the port number (0 relative)

and **iiiiiii** gives the individual id number of the adapter

so 00066a**00**a0iiiiiii is the port guid of the 1st port of the adapter

and 00066a**01**a0iiiiiii is the port guid of the 2nd port of the adapter.

Similarly, a VFx Port GUID is of the form **00066app38iiiiii**

where **pp** gives the IOC number (1 or 2)

and **iiiiiii** gives the individual ID number of the VIO hardware

so 00066a**01**38iiiiiii is the port guid of IOC 1 of VIO hardware iiiiiii

and 00066a**02**38iiiiiii is the port guid of IOC 2 of VIO hardware iiiiiii

NOTE

After a virtual adapter has been successfully added (meaning at least 1 session where part of the adapter has gone to the Active state) the SRP module will indicates what type of session was created in the mode variable (**ib_qlgc_srp_stats**) file, depending on whether "**roundrobinmode: 1**" is set in the **qlgc_srp.cfg** file. In this case "X" is the virtual adapter number, with number **0** being the first one created.

If no sessions were successfully brought to the Active state, then the **roundrobin_X** or **failover_X** file will not be created.

In a round robin configuration, if everything is configured correctly, all sessions will be Active.

In a failover configuration, if everything is configured correctly, one session will be Active and the rest will be Connected. The transition of a session from Connected to Active will not be attempted until that session needs to become Active, due to the failure of the previously Active session.

How does the user find a HCA port GUID?

Solution:

A HCA Port GUID is displayed by entering the following at any host prompt:

```
ibv_devinfo -i 1 for port 1
ibv_devinfo -i 2 for port 2
```

The system displays information similar to the following:

```
st106:~ # ibv_devinfo -i 1
hca_id: mthca0
fw_ver: 5.1.9301
node_guid: 0006:6a00:9800:6c9f
sys_image_guid: 0006:6a00:9800:6c9f
vendor_id: 0x066a
vendor_part_id: 25218
hw_ver: 0xA0
board_id: SS_0000000005
phys_port_cnt: 2
port:  1
state: PORT_ACTIVE (4)
max_mtu: 2048 (4)
active_mtu: 2048 (4)
sm_lid: 71
port_lid: 60
port_lmc: 0x00

st106:~ # ibv_devinfo -i 2
hca_id: mthca0
fw_ver: 5.1.9301
node_guid: 0006:6a00:9800:6c9f
sys_image_guid: 0006:6a00:9800:6c9f
vendor_id: 0x066a
vendor_part_id: 25218
hw_ver: 0xA0
board_id: SS_0000000005
phys_port_cnt: 2
port:  2
state: PORT_ACTIVE (4)
max_mtu: 2048 (4)
active_mtu: 2048 (4)
sm_lid: 71
port_lid: 64
port_lmc: 0x00
```

Need to determine the SRP driver version.

Solution:

To determine the SRP driver version number, enter the command **`modinfo -d qlgc-srp`**, which returns information similar to the following:

```
st159:~ # modinfo -d qlgc-srp
QLogic Corp. Virtual HBA (SRP) SCSI Driver, version 1.0.0.0.3
```

F Write Combining

Introduction

Write Combining improves write bandwidth to the QLogic driver by writing multiple words in a single bus transaction (typically 64 bytes). Write combining applies only to x86_64 systems.

The x86 Page Attribute Table (PAT) mechanism allocates Write Combining (WC) mappings for the PIO buffers, and is the default mechanism for WC.

If PAT is unavailable or PAT initialization fails, the software will generate a message in the log and fall back to the Memory Type Range Registers (MTRR) mechanism. If write combining is not working properly, lower than expected bandwidth may occur.

The following sections provide instructions for enabling and disabling WC using PAT and MTRR, and for verifying that write combining is working.

PAT and Write Combining

The `wc_pat` parameter is set in `/etc/modprobe.conf` (on Red Hat systems) or `/etc/modprobe.conf.local` (on SLES systems) to:

- 0) Disable PAT and use MTRR.
- 1) Configure WC by programming the PAT at the memory page level instead of the physical memory ranges.
- 2) Configure WC by programming the PAT at the memory page level and overwriting the operating system PAT configuration to enable WC uniformly across CPUs that have it disabled. This is the default setting.

The default `wc_pat` parameter is:

```
option ib_qib wc_pat=2
```

If PAT is unavailable or PAT initialization fails, the code generates a message in the log and falls back to the MTRR mechanism. To use MTRR, disable PAT by setting the `wc_pat` parameter to 0 (as a root user):

```
option ib_qib wc_pat=0
```

Revert to using MTRR-only behavior by following one of the two suggestions in MTRR Mapping and Write Combining.

The driver must be restarted after the changes have been made.

NOTE

There will not be a WC entry in `/proc/mtrr` when using PAT.

MTRR Mapping and Write Combining

Two suggestions for properly enabling MTRR mapping for write combining are described in the following sections.

See "Performance Issues" on page D-9 for more details on a related performance issue.

Edit BIOS Settings to Fix MTRR Issues

You can edit the BIOS setting for MTRR mapping. The BIOS setting looks similar to:

```
MTRR Mapping          [Discrete]
```

For systems with very large amounts of memory (32GB or more), it may also be necessary to adjust the BIOS setting for the *PCI hole granularity* to 2GB. This setting allows the memory to be mapped with fewer MTRRs, so that there will be one or more unused MTRRs for the InfiniPath driver.

Some BIOS' do not have the MTRR mapping option. It may have a different name, depending on the chipset, vendor, BIOS, or other factors. For example, it is sometimes referred to as *32 bit memory hole*. This setting must be enabled.

If there is no setting for MTRR mapping or 32 bit memory hole, and you have problems with degraded performance, contact your system or motherboard vendor and ask how to enable write combining.

Use the `ipath_mtrr` Script to Fix MTRR Issues

QLogic also provides a script, `ipath_mtrr`, which sets the MTRR registers, enabling maximum performance from the InfiniPath driver. This Python script is available as a part of the InfiniPath software download, and is contained in the `infinipath*` RPM. It is installed in `/bin`.

To diagnose the machine, run it with no arguments (as a root user):

```
# ipath_mtrr
```

The test results will list any problems, if they exist, and provide suggestions on what to do.

To fix the MTRR registers, use:

```
# ipath_mtrr -w
```

Restart the driver after fixing the registers.

This script needs to be run after each system reboot. It can be set to run automatically upon restart by adding this line in `/etc/sysconfig/infinipath`:

```
IPATH_MTRR_ACTIVE=1
```

See the `ipath_mtrr(8)` man page for more information on other options.

Verify Write Combining is Working

To see if write combining is working correctly and to check the bandwidth, run the following command:

```
$ ipath_pkt_test -B
```

With write combining enabled, the QLE7140 and QLE7240 report in the range of 1150–1500 MBps. The QLE7280 reports in the range of 1950–3000 MBps.

You can also use `ipath_checkout` (use option 5) to check bandwidth.

Increased latency and low bandwidth may indicate a problem. The interconnect could be operating in a degraded performance mode with latency increasing to several microseconds, and bandwidth decreasing to as little as 200 MBps.

Upon driver startup, you may see these errors:

```
ib_qib 0000:04:01.0: infinqib0: Performance problem: bandwidth
to PIO buffers is only 273 MiB/sec
.
.
.
```

If you do not see any of these messages on your console, but suspect this problem, check the `/var/log/messages` file. Some systems suppress driver load messages but still output them to the log file.

Notes

G Commands and Files

The most useful commands and files for debugging, and common tasks, are presented in the following sections. Many of these commands and files have been discussed elsewhere in the documentation. This information is summarized and repeated here for your convenience.

Check Cluster Homogeneity with `ipath_checkout`

Many problems can be attributed to the lack of homogeneity in the cluster environment. Use the following items as a checklist for verifying homogeneity. A difference in any one of these items in your cluster may cause problems:

- Kernels
- Distributions
- Versions of the QLogic boards
- Runtime and build environments
- `.o` files from different compilers
- Libraries
- Processor/link speeds
- PIO bandwidth
- MTUs

With the exception of finding any differences between the runtime and build environments, `ipath_checkout` will pick up information on all the above items. Other programs useful for verifying homogeneity are listed in Table G-1. More details on `ipath_checkout` are in “ipath_checkout” on page G-25.

Restarting InfiniPath

When the driver status appears abnormal on any node, you can try restarting (as a root user). Type:

```
# /etc/init.d/openibd restart
```

These two commands perform the same function as `restart`:

```
# /etc/init.d/openibd stop
# /etc/init.d/openibd start
```

Also check the `/var/log/messages` file for any abnormal activity.

Summary and Descriptions of Commands

Commands are summarized in Table G-1. Names in blue text are linked to a corresponding section that provides further details. Check the man pages for more information on the programs.

Table G-1. Useful Programs

Program Name	Function
`chkconfig`	Checks the configuration state and enables/disables services, including drivers. Can be useful for checking homogeneity.
`dmesg`	Prints out bootup messages. Useful for checking for initialization problems.
`iba_opp_query`	Retrieves path records from the Distributed SA and is somewhat similar to `iba_saquery`. It is intended for testing the Distributed SA (`qlogic_sa`) and for verifying connectivity between nodes in the fabric.
iba_hca_rev	Scans the system and reports hardware and firmware information about all the HCAs in the system.
iba_manage_switch	Allows management of externally managed switches (including 12200, 12200-18, and HP BLc QLogic 4X QDR) without the IFS software.
`iba_packet_capture`	Enables packet capture and subsequent dump to file
`ibhosts`[a]	Checks that all hosts in the fabric are up and visible to the subnet manager and to each other
`ibstatus`[a]	Checks the status of IB devices when OpenFabrics is installed

Table G-1. Useful Programs (Continued)

Program Name	Function
ibtracert[a]	Determines the path that IB packets travel between two nodes
ibv_devinfo[a]	Lists information about IB devices in use. Use when OpenFabrics is enabled.
ident[b]	Identifies RCS keyword strings in files. Can check for dates, release versions, and other identifying information.
ipath_checkout[c]	A bash shell script that performs sanity testing on a cluster using QLogic hardware and InfiniPath software. When the program runs without errors, the node is properly configured.
ipath_control[c]	A shell script that manipulates various parameters for the InfiniPath driver. This script gathers the same information contained in boardversion, status_str, and version.
ipath_mtrr[c]	A Python script that sets the MTRR registers.
ipath_pkt_test[c]	Tests the IB link and bandwidth between two QLogic IB adapters, or, using an IB loopback connector, tests within a single QLogic IB adapter
ipathstats[c]	Displays driver statistics and hardware counters, including performance and "error" (including status) counters
lsmod	Shows status of modules in the Linux kernel. Use to check whether drivers are loaded.
modprobe	Adds or removes modules from the Linux kernel.
mpirun[c]	A front end program that starts an MPI job on an InfiniPath cluster. Use to check the origin of the drivers.
mpi_stress	An MPI stress test program designed to load up an MPI interconnect with point-to-point messages while optionally checking for data integrity.
ps	Displays information on current active processes. Use to check whether all necessary processes have been started.
rpm	Package manager to install, query, verify, update, or erase software packages. Use to check the contents of a package.
strings[d]	Prints the strings of printable characters in a file. Useful for determining contents of non-text files such as date and version.

Table Notes

[a] These programs are contained in the OpenFabrics `openib-diags` RPM.

[b] These programs are contained within the `rcs` RPM for your distribution.

[c] These programs are contained in the Open `mpi-frontend` RPM.

[d] These programs are contained within the `binutils` RPM for your distribution.

dmesg

`dmesg` prints out bootup messages. It is useful for checking for initialization problems. You can check to see if problems were detected during the driver and QLogic hardware initialization with the command:

```
$ dmesg|egrep -i infinipath|qib
```

This command may generate more than one screen of output.

iba_opp_query

This command retrieves path records from the Distributed SA and is somewhat similar to `iba_saquery`. It is intended for testing the Distributed SA (`qlogic_sa`) and for verifying connectivity between nodes in the fabric. For information on configuring and using the Distributed SA, refer to "QLogic Distributed Subnet Administration" on page 3-12.

`iba_opp_query` does not access the SM when doing queries, it only accesses the local Distributed SA database. For that reason, the kinds of queries that can be done are much more limited than with `iba_saquery`. In particular, it can only find paths that start on the machine where the command is run. (In other words, the source LID or source GID must be on the local node.) In addition, queries must supply either a source and destination LID, or a source and destination GID. They cannot be mixed. In addition, you will usually need to provide either a SID that was specified in Distributed SA configuration file, or a pkey that matches such a SID.

Usage

```
iba_opp_query [-v level] [-hca hca] [-p port] [-s LID] [-d
LID] [-S GID] [-D GID] [-k pkey] [-i sid] [-H]
```

Options

`-v/--verbose` *`level`* — Debug level. Should be a number between 1 and 7. Default is 5.

`-s/--slid` ***`LID`*** — Source LID. Can be in decimal, hex (0x##) or octal (0##)

`-d/--dlid` ***`LID`*** — Destination LID. Can be in decimal, hex (0x##) or octal (0##)

-S/**--sgid** ***GID*** — Source GID. (Can be in GID (“0x########:0x########”) or inet6 format (“##:##:##:##:##:##:##:##”))

-D/**--dgid** ***GID*** — Destination GID. (Can be in GID (“0x########:0x########”) or inet6 format (“##:##:##:##:##:##:##:##”))

-k/**--pkey** ***pkey*** — Partition Key

-i/**--sid** ***sid*** — Service ID

-h/**--hca** ***hca*** — The HCA to use. (Defaults to the first HCA.) The HCA can be identified by name (“mthca0”, “qib1”, et cetera) or by number (1, 2, 3, et cetera).

-p/**--port** *port* — The port to use. (Defaults to the first port)

-H/**--help** — Provides this help text.

All arguments are optional, but ill-formed queries can be expected to fail. You must provide at least a pair of LIDs, or a pair of GIDs.

Sample output:

```
# iba_opp_query --slid 0x31 --dlid 0x75 --sid 0x107
Query Parameters:
    resv1                   0x0000000000000107
    dgid                    ::
    sgid                    ::
    dlid                    0x75
    slid                    0x31
    hop                     0x0
    flow                    0x0
    tclass                  0x0
    num_path                0x0
    pkey                    0x0
    qos_class               0x0
    sl                      0x0
    mtu                     0x0
    rate                    0x0
    pkt_life                0x0
    preference              0x0
    resv2                   0x0
    resv3                   0x0
Using HCA qib0
Result:
    resv1                   0x0000000000000107
    dgid                    fe80::11:7500:79:e54a
    sgid                    fe80::11:7500:79:e416
    dlid                    0x75
    slid                    0x31
    hop                     0x0
    flow                    0x0
    tclass                  0x0
    num_path                0x0
    pkey                    0xffff
    qos_class               0x0
    sl                      0x1
    mtu                     0x4
    rate                    0x6
    pkt_life                0x10
    preference              0x0
```

```
resv2                   0x0
resv3                   0x0
```

Explanation of Sample Output:

This is a simple query, specifying the source and destination LIDs and the desired SID. The first half of the output shows the full “query” that will be sent to the Distributed SA. Unused fields are set to zero or are blank.

In the center, the line “`Using HCA qib0`” tells us that, because we did not specify which HCA to query against, the tool chose one for us. (Normally, the user will never have to specify which HCA to use. This is only relevant in the case where a single node is connected to multiple physical IB fabrics.)

Finally, the bottom half of the output shows the result of the query. Note that, if the query had failed (because the destination does not exist or because the SID is not found in the Distributed SA) you will receive and error instead:

```
# iba_opp_query --slid 0x31 --dlid 0x75 --sid 0x108
Query Parameters:
    resv1                   0x0000000000000108
    dgid                    ::
    sgid                    ::
    dlid                    0x75
    slid                    0x31
    hop                     0x0
    flow                    0x0
    tclass                  0x0
    num_path                0x0
    pkey                    0x0
    qos_class               0x0
    sl                      0x0
    mtu                     0x0
    rate                    0x0
    pkt_life                0x0
    preference              0x0
    resv2                   0x0
    resv3                   0x0
Using HCA qib0
******
Error: Get Path returned 22 for query: Invalid argument
******
```

Examples:

Query by LID and SID:

```
iba_opp_query -s 0x31 -d 0x75 -i 0x107
iba_opp_query --slid 0x31 --dlid 0x75 --sid 0x107
```

Queries using octal or decimal numbers:

```
iba_opp_query --slid 061 --dlid 0165 --sid 0407 (using octal
numbers)
iba_opp_query –slid 49 –dlid 113 –sid 263 (using decimal
numbers)
```

Note that these queries are the same as the first two, only the base of the numbers has changed.

Query by LID and PKEY:

```
iba_opp_query --slid 0x31 --dlid 0x75 –pkey 0x8002
```

Query by GID:

```
iba_opp_query -S fe80::11:7500:79:e416 -D
fe80::11:7500:79:e54a --sid 0x107
iba_opp_query -S 0xfe80000000000000:0x001175000079e416 -D
0xfe80000000000000:0x001175000079e394 --sid 0x107
```

As before, these queries are identical to the first two queries – they are just using the GIDs instead of the LIDs to specify the ports involved.

iba_hca_rev

This command scans the system and reports hardware and firmware information about all the HCAs in the system. Running `iba_hca_rev -v`(as a root user) produces output similar to the following when run from a node on the IB fabric:

```
# iba_hca_rev -v
#####################
st2092  - HCA 0a:00.0
ID: FALCON QDR
PN: MHQH29B-XTR
EC: A2
SN: MT1029X00540
V0: PCIe Gen2 x8
V1: N/A
YA: N/A
FW: 2.9.1000
Image type:      ConnectX
FW Version:      2.9.1000
Device ID:       26428
Description:     Node             Port1            Port2            Sys image
GUIDs:           0002c903000ba8e0 0002c903000ba8e1 0002c903000ba8e2
0002c903000ba8e3
MACs:                                 0002c90ba8e0     0002c90ba8e1
Board ID:         (MT_0D80120009)
VSD:
PSID:            MT_0D80120009

Firmware Configuration:
;; Generated automatically by iniprep tool on Sun Jun 05 11:50:37 IDT 2011 from
./b0_falcon.prs
;;
;; PRS  FILE FOR Eagle
;; $Id: b0_falcon.prs,v 1.18 2011-02-14 11:47:28 achiad Exp $


[PS_INFO]
Name = MHQH29B-XTR_A2
Description = ConnectX-2 VPI adapter card; dual-port; 40Gb/s QSFP; PCIe2.0 x8
5.0GT/s; tall bracket; RoHS R6
```

```
[ADAPTER]
PSID = MT_0D80120009
pcie_gen2_speed_supported = true
adapter_dev_id = 0x673c
silicon_rev = 0xb0

gpio_mode1 = 0x0
gpio_mode0 = 0x050e070f
gpio_default_val = 0x0502010f

[HCA]
hca_header_device_id = 0x673c
hca_header_subsystem_id = 0x0017
dpdp_en = true
eth_xfi_en = true
mdio_en_port1 = 0
[IB]
phy_type_port1 = XFI
phy_type_port2 = XFI

read_cable_params_port1_en = true
read_cable_params_port2_en = true

;;Polarity
eth_tx_lane_polarity_port1=0x0
eth_tx_lane_polarity_port2=0x0
eth_rx_lane_polarity_port1=0x1
eth_rx_lane_polarity_port2=0xD

;;Lane reversal
eth_tx_lane_reversal_port1=on
eth_tx_lane_reversal_port2=on
eth_rx_lane_reversal_port1=on
eth_rx_lane_reversal_port2=on

port1_sd0_ob_preemp_pre_qdr = 0x0
port2_sd0_ob_preemp_pre_qdr = 0x0
port1_sd1_ob_preemp_pre_qdr = 0x0
port2_sd1_ob_preemp_pre_qdr = 0x0
```

```
port1_sd2_ob_preemp_pre_qdr = 0x0
port2_sd2_ob_preemp_pre_qdr = 0x0
port1_sd3_ob_preemp_pre_qdr = 0x0
port2_sd3_ob_preemp_pre_qdr = 0x0

port1_sd0_ob_preemp_post_qdr = 0x6
port2_sd0_ob_preemp_post_qdr = 0x6
port1_sd1_ob_preemp_post_qdr = 0x6
port2_sd1_ob_preemp_post_qdr = 0x6
port1_sd2_ob_preemp_post_qdr = 0x6
port2_sd2_ob_preemp_post_qdr = 0x6
port1_sd3_ob_preemp_post_qdr = 0x6
port2_sd3_ob_preemp_post_qdr = 0x6

port1_sd0_ob_preemp_main_qdr = 0x0
port2_sd0_ob_preemp_main_qdr = 0x0
port1_sd1_ob_preemp_main_qdr = 0x0
port2_sd1_ob_preemp_main_qdr = 0x0
port1_sd2_ob_preemp_main_qdr = 0x0
port2_sd2_ob_preemp_main_qdr = 0x0
port1_sd3_ob_preemp_main_qdr = 0x0
port2_sd3_ob_preemp_main_qdr = 0x0

port1_sd0_ob_preemp_msb_qdr = 0x0
port2_sd0_ob_preemp_msb_qdr = 0x0
port1_sd1_ob_preemp_msb_qdr = 0x0
port2_sd1_ob_preemp_msb_qdr = 0x0
port1_sd2_ob_preemp_msb_qdr = 0x0
port2_sd2_ob_preemp_msb_qdr = 0x0
port1_sd3_ob_preemp_msb_qdr = 0x0
port2_sd3_ob_preemp_msb_qdr = 0x0

port1_sd0_muxmain_qdr = 0x1f
port2_sd0_muxmain_qdr = 0x1f
port1_sd1_muxmain_qdr = 0x1f
port2_sd1_muxmain_qdr = 0x1f
port1_sd2_muxmain_qdr = 0x1f
port2_sd2_muxmain_qdr = 0x1f
port1_sd3_muxmain_qdr = 0x1f
```

```
port2_sd3_muxmain_qdr = 0x1f

mellanox_qdr_ib_support = true
mellanox_ddr_ib_support = true

spec1_2_ib_support = true
spec1_2_ddr_ib_support = true
spec1_2_qdr_ib_support = true

auto_qdr_tx_options = 8
auto_qdr_rx_options = 7

auto_ddr_option_0.tx_preemp_pre = 0x2
auto_ddr_option_0.tx_preemp_msb = 0x1
auto_ddr_option_0.tx_preemp_post = 0x0
auto_ddr_option_0.tx_preemp_main = 0x1b

auto_ddr_option_1.tx_preemp_pre = 0x8
auto_ddr_option_1.tx_preemp_msb = 0x0
auto_ddr_option_1.tx_preemp_post = 0x2
auto_ddr_option_1.tx_preemp_main = 0x10
auto_ddr_option_1.tx_preemp = 0x0

auto_ddr_option_2.tx_preemp_pre = 0xa
auto_ddr_option_2.tx_preemp_msb = 0x0
auto_ddr_option_2.tx_preemp_post = 0x2
auto_ddr_option_2.tx_preemp_main = 0x12
auto_ddr_option_2.tx_preemp = 0x0

auto_ddr_option_3.tx_preemp_pre = 0xf
auto_ddr_option_3.tx_preemp_msb = 0x1
auto_ddr_option_3.tx_preemp_post = 0x3
auto_ddr_option_3.tx_preemp_main = 0x1f
auto_ddr_option_3.tx_preemp = 0x2

auto_ddr_option_4.tx_preemp_pre = 0x4
auto_ddr_option_4.tx_preemp_msb = 0x1
auto_ddr_option_4.tx_preemp_post = 0x5
auto_ddr_option_4.tx_preemp_main = 0x12
```

```
auto_ddr_option_4.tx_preemp = 0x0

auto_ddr_option_5.tx_preemp_pre = 0x5
auto_ddr_option_5.tx_preemp_msb = 0x1
auto_ddr_option_5.tx_preemp_post = 0x3
auto_ddr_option_5.tx_preemp_main = 0x13
auto_ddr_option_5.tx_preemp = 0x0

auto_ddr_option_6.tx_preemp_pre = 0x3
auto_ddr_option_6.tx_preemp_msb = 0x1
auto_ddr_option_6.tx_preemp_post = 0x4
auto_ddr_option_6.tx_preemp_main = 0x1f
auto_ddr_option_6.tx_preemp = 0x0

auto_ddr_option_7.tx_preemp_pre = 0x8
auto_ddr_option_7.tx_preemp_msb = 0x1
auto_ddr_option_7.tx_preemp_post = 0x3
auto_ddr_option_7.tx_preemp_main = 0x17
auto_ddr_option_7.tx_preemp = 0x0

auto_ddr_option_8.tx_preemp_pre = 0xf
auto_ddr_option_8.tx_preemp_msb = 0x1
auto_ddr_option_8.tx_preemp_post = 0x3
auto_ddr_option_8.tx_preemp_main = 0x14
auto_ddr_option_8.tx_preemp = 0x2

auto_ddr_option_9.tx_preemp_pre = 0x8
auto_ddr_option_9.tx_preemp_msb = 0x0
auto_ddr_option_9.tx_preemp_post = 0x3
auto_ddr_option_9.tx_preemp_main = 0x17
auto_ddr_option_9.tx_preemp = 0x0

auto_ddr_option_10.tx_preemp_pre = 0x8
auto_ddr_option_10.tx_preemp_msb = 0x0
auto_ddr_option_10.tx_preemp_post = 0x3
auto_ddr_option_10.tx_preemp_main = 0x17
auto_ddr_option_10.tx_preemp = 0x0

auto_ddr_option_11.tx_preemp_pre = 0xf
```

```
auto_ddr_option_11.tx_preemp_msb = 0x0
auto_ddr_option_11.tx_preemp_post = 0x3
auto_ddr_option_11.tx_preemp_main = 0x19
auto_ddr_option_11.tx_preemp = 0x0

auto_ddr_option_12.tx_preemp_pre = 0xf
auto_ddr_option_12.tx_preemp_msb = 0x0
auto_ddr_option_12.tx_preemp_post = 0x3
auto_ddr_option_12.tx_preemp_main = 0x19
auto_ddr_option_12.tx_preemp = 0x0

auto_ddr_option_13.tx_preemp_pre = 0x0
auto_ddr_option_13.tx_preemp_msb = 0x0
auto_ddr_option_13.tx_preemp_post = 0x0
auto_ddr_option_13.tx_preemp_main = 0x5
auto_ddr_option_13.tx_preemp = 0x0

auto_ddr_option_14.tx_preemp_pre = 0x0
auto_ddr_option_14.tx_preemp_msb = 0x0
auto_ddr_option_14.tx_preemp_post = 0x0
auto_ddr_option_14.tx_preemp_main = 0x5
auto_ddr_option_14.tx_preemp = 0x0

auto_ddr_option_15.tx_preemp_pre = 0x0
auto_ddr_option_15.tx_preemp_msb = 0x0
auto_ddr_option_15.tx_preemp_post = 0x0
auto_ddr_option_15.tx_preemp_main = 0x5
auto_ddr_option_15.tx_preemp = 0x0

;;;;; Integer parameter. Values range : 0x0 - 0xf.



auto_ddr_option_0.rx_offs_lowpass_en = 0x0
auto_ddr_option_1.rx_offs_lowpass_en = 0x0
auto_ddr_option_2.rx_offs_lowpass_en = 0x0
auto_ddr_option_3.rx_offs_lowpass_en = 0x0
auto_ddr_option_4.rx_offs_lowpass_en = 0x0
auto_ddr_option_5.rx_offs_lowpass_en = 0x0
```

```
auto_ddr_option_6.rx_offs_lowpass_en = 0x0
auto_ddr_option_7.rx_offs_lowpass_en = 0x0
auto_ddr_option_0.rx_offs = 0x0
auto_ddr_option_1.rx_offs = 0x0
auto_ddr_option_2.rx_offs = 0x0
auto_ddr_option_3.rx_offs = 0x0
auto_ddr_option_4.rx_offs = 0x0
auto_ddr_option_5.rx_offs = 0x0
auto_ddr_option_6.rx_offs = 0x0
auto_ddr_option_7.rx_offs = 0x0
auto_ddr_option_0.rx_equal_offs = 0x0
auto_ddr_option_1.rx_equal_offs = 0x0
auto_ddr_option_2.rx_equal_offs = 0x0
auto_ddr_option_3.rx_equal_offs = 0x0
auto_ddr_option_4.rx_equal_offs = 0x0
auto_ddr_option_5.rx_equal_offs = 0x0
auto_ddr_option_6.rx_equal_offs = 0x0
auto_ddr_option_7.rx_equal_offs = 0x0
auto_ddr_option_0.rx_muxeq = 0x0
auto_ddr_option_1.rx_muxeq = 0x0
auto_ddr_option_2.rx_muxeq = 0x0
auto_ddr_option_3.rx_muxeq = 0x0
auto_ddr_option_4.rx_muxeq = 0x0
auto_ddr_option_5.rx_muxeq = 0x0
auto_ddr_option_6.rx_muxeq = 0x0
auto_ddr_option_7.rx_muxeq = 0x0
auto_ddr_option_0.rx_muxmain = 0x1f
auto_ddr_option_1.rx_muxmain = 0x1f
auto_ddr_option_2.rx_muxmain = 0x1f
auto_ddr_option_3.rx_muxmain = 0x1f
auto_ddr_option_4.rx_muxmain = 0x1f
auto_ddr_option_5.rx_muxmain = 0x1f
auto_ddr_option_6.rx_muxmain = 0x1f
auto_ddr_option_7.rx_muxmain = 0x1f
auto_ddr_option_0.rx_main = 0x1
auto_ddr_option_1.rx_main = 0xf
auto_ddr_option_2.rx_main = 0xf
auto_ddr_option_3.rx_main = 0xf
auto_ddr_option_4.rx_main = 0xe
```

```
auto_ddr_option_5.rx_main = 0xe
auto_ddr_option_6.rx_main = 0xf
auto_ddr_option_7.rx_main = 0xf
auto_ddr_option_0.rx_extra_hs_gain = 0x0
auto_ddr_option_1.rx_extra_hs_gain = 0x3
auto_ddr_option_2.rx_extra_hs_gain = 0x2
auto_ddr_option_3.rx_extra_hs_gain = 0x4
auto_ddr_option_4.rx_extra_hs_gain = 0x1
auto_ddr_option_5.rx_extra_hs_gain = 0x2
auto_ddr_option_6.rx_extra_hs_gain = 0x7
auto_ddr_option_7.rx_extra_hs_gain = 0x0
auto_ddr_option_0.rx_sigdet_th = 0x1
auto_ddr_option_1.rx_sigdet_th = 0x1
auto_ddr_option_2.rx_sigdet_th = 0x1
auto_ddr_option_3.rx_sigdet_th = 0x1
auto_ddr_option_4.rx_sigdet_th = 0x1
auto_ddr_option_5.rx_sigdet_th = 0x1
auto_ddr_option_6.rx_sigdet_th = 0x1
auto_ddr_option_7.rx_sigdet_th = 0x1
auto_ddr_option_0.rx_equalization = 0x4
auto_ddr_option_1.rx_equalization = 0x0
auto_ddr_option_2.rx_equalization = 0x0
auto_ddr_option_3.rx_equalization = 0x0
auto_ddr_option_4.rx_equalization = 0x0
auto_ddr_option_5.rx_equalization = 0x0
auto_ddr_option_6.rx_equalization = 0x0
auto_ddr_option_7.rx_equalization = 0x0

auto_ddr_option_9.rx_muxeq = 0x0
auto_ddr_option_9.rx_muxmain = 0x1f
auto_ddr_option_9.rx_main = 0xf
auto_ddr_option_9.rx_extra_hs_gain = 0x0
auto_ddr_option_9.rx_equalization = 0x0

auto_ddr_option_10.rx_muxeq = 0x0
auto_ddr_option_10.rx_muxmain = 0x1f
auto_ddr_option_10.rx_main = 0xf
auto_ddr_option_10.rx_extra_hs_gain = 0x0
auto_ddr_option_10.rx_equalization = 0x0
```

```
auto_ddr_option_11.rx_muxeq = 0x04
auto_ddr_option_11.rx_muxmain = 0x1f
auto_ddr_option_11.rx_main = 0xf
auto_ddr_option_11.rx_extra_hs_gain = 0x4
auto_ddr_option_11.rx_equalization = 0x7f

auto_ddr_option_12.rx_muxeq = 0x6
auto_ddr_option_12.rx_muxmain = 0x1f
auto_ddr_option_12.rx_main = 0xf
auto_ddr_option_12.rx_extra_hs_gain = 0x4
auto_ddr_option_12.rx_equalization = 0x7f

auto_ddr_option_13.rx_muxeq = 0x0
auto_ddr_option_13.rx_muxmain = 0x1f
auto_ddr_option_13.rx_main = 0xf
auto_ddr_option_13.rx_extra_hs_gain = 0x3
auto_ddr_option_13.rx_equalization = 0x0

auto_ddr_option_14.rx_muxeq = 0x0
auto_ddr_option_14.rx_muxmain = 0x1f
auto_ddr_option_14.rx_main = 0xf
auto_ddr_option_14.rx_extra_hs_gain = 0x3
auto_ddr_option_14.rx_equalization = 0x0

auto_ddr_option_15.rx_muxeq = 0x0
auto_ddr_option_15.rx_muxmain = 0x1f
auto_ddr_option_15.rx_main = 0xf
auto_ddr_option_15.rx_extra_hs_gain = 0x3
auto_ddr_option_15.rx_equalization = 0x0

center_mix90phase = true
auto_kr_option_6.rx_extra_hs_gain = 0x3

ext_phy_board_port1 = FALCON
ext_phy_board_port2 = FALCON

[PLL]
lbist_en  = 0
```

```
lbist_shift_freq  = 3
pll_stabilize = 0x13
flash_div = 0x3
lbist_array_bypass = 1
lbist_pat_cnt_lsb = 0x2
core_f = 44
core_r = 27
ddr_6_db_preemp_pre = 0x3
ddr_6_db_preemp_main = 0xe

[FW]


Firmware Verification:

     FS2 failsafe image. Start address: 0x0. Chunk size 0x80000:

     NOTE: The addresses below are contiguous logical addresses. Physical
addresses on
           flash may be different, based on the image start address and chunk
size

     /0x00000038-0x00001233 (0x0011fc)/ (BOOT2) - OK
     /0x00001234-0x0000280f (0x0015dc)/ (BOOT2) - OK
     /0x00002810-0x000034ef (0x000ce0)/ (Configuration) - OK
     /0x000034f0-0x00003533 (0x000044)/ (GUID) - OK
     /0x00003534-0x0000366b (0x000138)/ (Image Info) - OK
     /0x0000366c-0x0000946f (0x005e04)/ (DDR) - OK
     /0x00009470-0x0000ab53 (0x0016e4)/ (DDR) - OK
     /0x0000ab54-0x00016b43 (0x00bff0)/ (DDR) - OK
     /0x00016b44-0x0001fb57 (0x009014)/ (DDR) - OK
     /0x0001fb58-0x000720ab (0x052554)/ (DDR) - OK
     /0x000720ac-0x0007308f (0x000fe4)/ (DDR) - OK
     /0x00073090-0x00099787 (0x0266f8)/ (DDR) - OK
     /0x00099788-0x0009d11f (0x003998)/ (DDR) - OK
     /0x0009d120-0x000a0b8b (0x003a6c)/ (DDR) - OK
     /0x000a0b8c-0x000a1037 (0x0004ac)/ (Configuration) - OK
     /0x000a1038-0x000a1093 (0x00005c)/ (Jump addresses) - OK
     /0x000a1094-0x000a1707 (0x000674)/ (FW Configuration) - OK
     /0x00000000-0x000a1707 (0x0a1708)/ (Full Image) - OK
```

```
FW image verification succeeded. Image is bootable.

######################
```

iba_manage_switch

(**Switch**) Allows management of externally managed switches (including 12200, 12200-18, and HP BLc QLogic 4X QDR) without using the IFS software. It is designed to operate on one switch at a time, taking a mandatory target GUID parameter.

Usage

```
iba_manage_switch -t target-guid [-H] [-v] [-h hca] [-p port]
[ -x] [-S] [ -f fileName] [-r] [-C configOption] [-i
integer-value] [-s string-value] [-c captureFile] operation
```

Options

-H – help (this message)

-v – verbose - additional output

-t *target-guid* – guid of target switch in hex format, for example `0x00066a00e3001234`

-h *hca* – HCA number, default is first HCA

-p *port* – port number, default is first active port

-x – clobber previous results file

-S – enforce password, will be prompted for each subcommand

-f *fileName* – fileName of the `emfw` file to be used in `fwUpdate` operation - must be a valid `emfw` file with `.emfw` suffix

-r – reset switch after `fwUpdate` (only valid with `fwUpdate` operation)

-C *configOption* – configuration option for `setConfigValue` operation

mtucap (mtu capability) – use `-i` for integer value (4-2048, 5-4096)

vlcap (vl capability) – use `-i` for integer value (1=1VL, 2=2VLs, 3=4VLs, 4=8VLs, 5=15VLs)

linkwidth (link width supported) – use `-i` for integer value (1=1X, 2=4X, 3=1X/4X, 4=8X, 5=1X/8X, 6=4X/8X, 7=1X/4X/8X)

vlcreditdist (VL credit distribution) – use `-i` for integer value (0, 1, 2, 3, or 4)

linkspeed (link speed supported) – use `-i` for integer value (1=SDR, 2=DDR, 3=SDR/DDR, 4=QDR, 7=SDR/DDR/QDR)

`-i` *integer-value* – integer value

-s *string-value* – string value

-c *captureFile* – filename of capture output file

operation – operation to perform:

fwUpdate – perform firmware update using `fileName` parameter, must be an `emfw` file

fwVerify – perform firmware validation, validate firmware in primary/secondary EEPROMs, report which was booted

ping – test for switch presence

reboot – reboot switch

setConfigValue – update configuration value, use `-C` for configuration option and `-i` for integer value

setIBNodeDesc – set the IB node description, use `-s` for string value of node desc

setPassword – set the vendor key (prompts for password to be used for subsequent switch access)

showConfig – report user-configurable settings

showFwVersion – report firmware version running on switch

showPowerCooling – report status of power supplies and fans

capture – perform capture of switch

showVPD – report VPD information of switch

Example

```
iba_manage_switch -t 0x00066a00e3001234 -f
  QLogic_12000_V1_firmware.7.0.0.0.27.emfw fwUpdate
iba_manage_switch -t 0x00066a00e3001234 reboot
iba_manage_switch -t 0x00066a00e3001234 showFwVersion
iba_manage_switch -t 0x00066a00e3001234 -s i12k1234
  setIBNodeDesc
iba_manage_switch -t 0x00066a00e3001234 -C mtucap -i 4
  setConfigValue
iba_manage_switch -H
```

The results are recorded in `iba_manage_switch.res` file in the current directory. Use the `-x` option to clobber and create a new file.

iba_packet_capture

This tool operates in cooperation with IB snoop device in the QIB driver. It enables packet capture and subsequent dump to file, the format of which is compatible with the Wireshark tool. Wireshark provides IB packet analysis.

The `snoop_enable` variable must be set to 1 (enabled) in the `modprobe.conf /ib_qib.conf` file to create snoop devices and capture devices. If `snoop_enable` is set to 0 (disable) then no snoop and capture devices are created.

This tool captures packets in memory in a large ring buffer and dumps the packet information to a file when it is instructed by a set option. The tool supports filtering of several IB fields. The tool is primarily intended for internal Intel use.

Edit the file `/etc/modprobe.d/ib_qib.conf` to add `snoop_enable=1` to the options line. If the file does not exist, it can be created with the following text:

options ib_qib snoop_enable=1

Usage

`iba_packet_capture [-o` *`outfile`*`] [-d` *`devfile`*`] [-f` *`filterfile`*`] [-a` *`alarm`*`] [-s` *`maxblocks`*`] [-v]`

Options

`-o` *`outfile`* – output file for captured packets - default is packetDump.pcap

`-d` *`devfile`* – snoop device file for capturing packets - default is /dev/ipath_capture_00_01

`-f` *`filterfile`* – filter file used for filtering - if absent, no filtering

`-a` *`alarm`* – number of seconds for alarm trigger to dump capture and exit

`-s` *`maxblocks`* – max 64 byte blocks of data to capture in units of Mi (1024*1024)

`-v` – verbose output

To stop capture and trigger dump, kill with SIGINT (Ctrl-C) or SIGUSR1 (with the kill command). The program will dump packets to file and exit

A sample filter file is located at `/opt/iba/samples/filterFile.txt`. This file should be copied to the user's home directory for editing and used with the packet capture utility.

Wireshark can be downloaded from http://www.wireshark.org/. Intel recommends using version 1.6.2.

ibhosts

This tool determines if all the hosts in your IB fabric are up and visible to the subnet manager and to each other. It is installed from the `openib-diag` RPM. Running `ibhosts` (as a root user) produces output similar to the following when run from a node on the IB fabric:

```
# ibhosts
Ca : 0x0008f10001280000 ports 2 "Voltaire InfiniBand
Fiber-Channel Router"
Ca : 0x0011750000ff9869 ports 1 "idev-11"
Ca : 0x0011750000ff9878 ports 1 "idev-05"
Ca : 0x0011750000ff985c ports 1 "idev-06"
Ca : 0x0011750000ff9873 ports 1 "idev-04"
```

ibstatus

This program displays basic information on the status of IB devices that are currently in use when OpenFabrics RPMs are installed. It is installed from the `openib-diag` RPM.

Following is a sample output for the DDR adapters:

```
# ibstatus
Infiniband device 'qib0' port 1 status:
        default gid:
fe80:0000:0000:0000:0011:7500:0078:a5d2
        base lid:        0x1
        sm lid:          0x4
        state:           4: ACTIVE
        phys state:      5: LinkUp
        rate:            40 Gb/sec (4X QDR)
        link_layer:      InfiniBand
```

ibtracert

The tool `ibtracert` determines the path that IB packets travel between two nodes. It is installed from the `openib-diag` RPM. The IB LIDs of the two nodes in this example are determined by using the `ipath_control -i` command on each node. The `ibtracert` tool produces output similar to the following when run (as a root user) from a node on the IB fabric:

```
# ibtracert 0xb9 0x9a
   From ca {0x0011750000ff9886} portnum 1 lid 0xb9-0xb9
"iqa-37"
   [1] -> switch port {0x0002c9010a19bea0}[1] lid 0x14-0x14
"MT47396 Infiniscale-III"
   [24] -> switch port {0x00066a0007000333}[8] lid 0xc-0xc
"SilverStorm 9120 GUID=0x00066a000200016c Leaf 6, Chip A"
   [6] -> switch port {0x0002c90000000000}[15] lid 0x9-0x9
"MT47396 Infiniscale-III"
   [7] -> ca port {0x0011750000ff9878}[1] lid 0x9a-0x9a
"idev-05"
   To ca {0x0011750000ff9878} portnum 1 lid 0x9a-0x9a
"idev-05"
```

ibv_devinfo

This program displays information about IB devices, including various kinds of identification and status data. It is installed from the `openib-diag` RPM. Use this program when OpenFabrics is enabled. `ibv_devinfo` queries RDMA devices. Use the `-v` option to see more information. For example:

```
# ibv_devinfo
hca_id: qib0
        transport:                      InfiniBand (0)
        fw_ver:                         0.0.0
        node_guid:                      0011:7500:0078:a5d2
        sys_image_guid:                 0011:7500:0078:a5d2
        vendor_id:                      0x1175
        vendor_part_id:                 29474
        hw_ver:                         0x2
        board_id:                       InfiniPath_QLE7340
        phys_port_cnt:                  1
                port:   1
                       state:                   PORT_ACTIVE (4)
                        max_mtu:                 4096 (5)
                        active_mtu:              4096 (5)
                        sm_lid:                  4
                        port_lid:                1
                        port_lmc:                0x00
                        link_layer:              IB
```

ident

The `ident` strings are available in `ib_qib.ko`. Running `ident` provides driver information similar to the following. For QLogic RPMs on a SLES distribution, it will look like the following example:

ident/lib/modules/*OS_version*/updates/kernel/drivers/infiniband/hw/qib/ib_qib.ko

```
/lib/modules/OS_version/updates/kernel/drivers/infiniband/hw/
qib/ib_qib.ko:
$Id: QLogic OFED Release x.x.x $
$Date: yyyy-mm-dd-hh:mm $
```

NOTE

For QLogic RPMs on a RHEL distribution, the drivers folder is in the updates folder instead of the kernels folder as follows:

```
/lib/modules/OS_version/updates/drivers/
infiniband/hw/qib/ib_qib.ko
```

If the `/lib/modules/OS_version/updates` directory is not present, then the driver in use is the one that comes with the core kernel. In this case, either the `kernel-ib` RPM is not installed or it is not configured for the current running kernel.

If the `updates` directory is present, but empty except for the subdirectory `kernel`, then an OFED install is probably being used, and the `ident` string will be empty. For example:

```
$ cd /lib/modules/OS_version/updates
$ ls
kernel
$ cd kernel/drivers/infiniband/hw/qib/
lib/modules/2.6.18-8.el5/updates/kernel/drivers/infiniband/hw
/qib
$ ident ib_qib.ko
        ib_qib.ko:
        ident warning: no id keywords in ib_qib.ko
```

NOTE

`ident` is in the optional `rcs` RPM, and is not always installed.

ipath_checkout

The `ipath_checkout` tool is a `bash` script that verifies that the installation is correct and that all the nodes of the network are functioning and mutually connected by the InfiniPath fabric. It is installed from the `infinipath` RPM. It must be run on a front end node, and requires specification of a nodefile. For example:

```
$ ipath_checkout [options] nodefile
```

The `nodefile` lists the hostnames of the nodes of the cluster, one hostname per line. The format of `nodefile` is as follows:

```
hostname1
hostname2
...
```

NOTE

- The hostnames in the `nodefile` are Ethernet hostnames, not IPv4 addresses.
- To create a `nodefile`, use the `ibhosts` program. It will generate a list of available nodes that are already connected to the switch.

`ipath_checkout` performs the following seven tests on the cluster:

1. Executes the `ping` command to all nodes to verify that they all are reachable from the front end.
2. Executes the `ssh` command to each node to verify correct configuration of `ssh`.
3. Gathers and analyzes system configuration from the nodes.
4. Gathers and analyzes RPMs installed on the nodes.
5. Verifies InfiniPath hardware and software status and configuration, including tests for link speed, PIO bandwidth (incorrect MTRR settings), and MTU size.
6. Verifies the ability to `mpirun` jobs on the nodes.
7. Runs a bandwidth and latency test on every pair of nodes and analyzes the results.

The options available with `ipath_checkout` are shown in Table G-2.

Table G-2. `ipath_checkout` Options

Command	Meaning
`-h, --help`	These options display help messages describing how a command is used.

`-v, --verbose`
`-vv, --vverbose`
`-vvv, --vvverbose`
	These options specify three successively higher levels of detail in reporting test results. There are four levels of detail in all, including the case where none of these options are given.

`-c, --continue`	When this option is not specified, the test terminates when any test fails. When specified, the tests continue after a failure, with failing nodes excluded from subsequent tests.

Table G-2. `ipath_checkout` Options (Continued)

Command	Meaning
`-k, --keep`	This option keeps intermediate files that were created while performing tests and compiling reports. Results are saved in a directory created by `mktemp` and named `infinipath_XXXXXX` or in the directory name given to `--workdir`.
`--workdir=DIR`	Use `DIR` to hold intermediate files created while running tests. `DIR` must not already exist.
`--run=LIST`	This option runs only the tests in `LIST`. See the seven tests listed previously. For example, `--run=123` will run only tests 1, 2, and 3.
`--skip=LIST`	This option skips the tests in `LIST`. See the seven tests listed previously. For example, `--skip=2457` will skip tests 2, 4, 5, and 7.
`-d, --debug`	This option turns on the `-x` and `-v` flags in `bash(1)`.

In most cases of failure, the script suggests recommended actions. Also refer to the `ipath_checkout` man page.

ipath_control

The `ipath_control` tool is a shell script that manipulates various parameters for the InfiniPath driver. It is installed from the `infinipath` RPM. Many of the parameters are used only when diagnosing problems, and may require special system configurations. Using these options may require restarting the driver or utility programs to recover from incorrect parameters.

Most of the functionality is accessed via the `/sys` filesystem. This shell script gathers the same information contained in these files:

```
/sys/class/infiniband/qib0/device/boardversion
/sys/class/infiniband/qib0/ports/1/linkcontrol/status_str
/sys/class/infiniband/qib0/device/driver/version
```

These files are also documented in Table G-4 and Table G-5.

Other than the `-i` option, this script must be run with root permissions. See the man pages for `ipath_control` for more details.

Here is sample usage and output:

```
% ipath_control -i
$Id: QLogic OFED Release x.x.x $ $Date: yyyy-mm-dd-hh:mm $
0: Version: ChipABI 2.0, InfiniPath_QLE7342, InfiniPath1 6.1,
SW Compat 2
0: Serial: RIB0941C00005 LocalBus: PCIe,5000MHz,x8
0,1: Status: 0xe1 Initted Present IB_link_up IB_configured
0,1: LID=0x1 GUID=0011:7500:0079:e574
0,1: HRTBT:Auto LINK:40 Gb/sec (4X QDR)
0,2: Status: 0x21 Initted Present [IB link not Active]
0,2: LID=0xffff GUID=0011:7500:0079:e575
```

The `-i` option combined with the `-v` option is very useful for looking at the IB width/rate and PCIe lanes/rate. For example:

```
% ipath_control -iv
$Id: QLogic OFED Release x.x.x $ $Date: yyyy-mm-dd-hh:mm $
0: Version: ChipABI 2.0, InfiniPath_QLE7342, InfiniPath1 6.1,
SW Compat 2
0: Serial: RIB0941C00005 LocalBus: PCIe,5000MHz,x8
0,1: Status: 0xe1 Initted Present IB_link_up IB_configured
0,1: LID=0x1 GUID=0011:7500:0079:e574
0,1: HRTBT:Auto LINK:40 Gb/sec (4X QDR)
0,2: Status: 0x21 Initted Present [IB link not Active]
0,2: LID=0xffff GUID=0011:7500:0079:e575
0,2: HRTBT:Auto LINK:10 Gb/sec (4X)
```

NOTE

On the first line, `Release` *`version`* refers to the current software release. The second line contains chip architecture version information.

Another useful option blinks the LED on the InfiniPath adapter (QLE7240 and QLE7280 adapters). This is useful for finding an adapter within a cluster. Run the following as a root user:

```
# ipath_control -b [On|Off]
```

ipath_mtrr

NOTE

Use `ipath_mtrr` if you are not using the default PAT mechanism to enable write combining.

MTRR is used by the InfiniPath driver to enable write combining to the QLogic on-chip transmit buffers. This option improves write bandwidth to the QLogic chip by writing multiple words in a single bus transaction (typically 64 bytes). This option applies only to x86_64 systems. It can often be set in the BIOS.

However, some BIOS' do not have the MTRR mapping option. It may have a different name, depending on the chipset, vendor, BIOS, or other factors. For example, it is sometimes referred to as *32 bit memory hole.* This setting must be enabled.

If there is no setting for MTRR mapping or 32 bit memory hole, contact your system or motherboard vendor and ask how to enable write combining.

You can check and adjust these BIOS settings using the BIOS Setup utility. For specific instructions, follow the hardware documentation that came with your system.

QLogic also provides a script, `ipath_mtrr`, which sets the MTRR registers, enabling maximum performance from the InfiniPath driver. This Python script is available as a part of the InfiniPath software download, and is contained in the `infinipath*` RPM. It is installed in `/bin`.

To diagnose the machine, run it with no arguments (as a root user):

```
# ipath_mtrr
```

The test results will list any problems, if they exist, and provide suggestions on what to do.

To fix the MTRR registers, use:

```
# ipath_mtrr -w
```

Restart the driver after fixing the registers.

This script needs to be run after each system reboot. It can be set to run automatically upon restart by adding this line in `/etc/sysconfig/infinipath`:

```
IPATH_MTRR_ACTIVE=1
```

See the `ipath_mtrr(8)` man page for more information on other options.

ipath_pkt_test

This program is installed from the `infinipath` RPM. Use `ipath_pkt_test` to do one of the following:

- Test the IB link and bandwidth between two InfiniPath IB adapters.
- Using an IB loopback connector, test the link and bandwidth within a single InfiniPath IB adapter.

The `ipath_pkt_test` program runs in either ping-pong mode (send a packet, wait for a reply, repeat) or in stream mode (send packets as quickly as possible, receive responses as they come back).

Upon completion, the sending side prints statistics on the packet bandwidth, showing both the payload bandwidth and the total bandwidth (including IB and InfiniPath headers). See the man page for more information.

ipathstats

The `ipathstats` program is useful for diagnosing InfiniPath problems, particularly those that are performance related. It is installed from the `infinipath` RPM. It displays both driver statistics and hardware counters, including both performance and "error" (including status) counters.

Running `ipathstats -c 10`, for example, displays the number of packets and 32-bit words of data being transferred on a node in each 10-second interval. This output may show differences in traffic patterns on different nodes, or at different stages of execution. See the man page for more information.

lsmod

When you need to find which InfiniPath and OpenFabrics modules are running, type the following command:

```
# lsmod | egrep 'ib_|rdma_|findex'
```

modprobe

Use this program to load/unload the drivers. You can check to see if the driver has loaded by using this command:

```
# modprobe -v ib_qib
```

The `-v` option typically only prints messages if there are problems.

The configuration file that `modprobe` uses is `/etc/modprobe.conf` (`/etc/modprobe.conf.local` on SLES). In this file, various options and naming aliases can be set.

mpirun

`mpirun` determines whether the program is being run against a QLogic or non-QLogic driver. It is installed from the `mpi-frontend` RPM. Sample commands and results are shown in the following paragraphs.

QLogic-built:

```
$ mpirun -np 2 -m /tmp/id1 -d0x101 mpi_latency 1 0
asus-01:0.ipath_setaffinity: Set CPU affinity to 1, port 0:2:0
(1 active chips)
asus-01:0.ipath_userinit: Driver is QLogic-built
```

Non-QLogic built:

```
$ mpirun -np 2 -m /tmp/id1 -d0x101 mpi_latency 1 0
asus-01:0.ipath_setaffinity: Set CPU affinity to 1, port 0:2:0
(1 active chips)
asus-01:0.ipath_userinit: Driver is not QLogic-built
```

mpi_stress

This is an MPI stress test program designed to load up an MPI interconnect with point-to-point messages while optionally checking for data integrity. By default, it runs with all-to-all traffic patterns, optionally including oneself and one's local shared memory (shm) peers. It can also be set up with multi-dimensional grid traffic patterns; this can be parameterized to run rings, open 2D grids, closed 2D grids, cubic lattices, hypercubes, and so on.

Optionally, the message data can be randomized and checked using CRC checksums (strong but slow) or XOR checksums (weak but fast). The communication kernel is built out of non-blocking point-to-point calls to load up the interconnect. The program is not designed to exhaustively test out different MPI primitives. Performance metrics are displayed, but should be carefully interpreted in terms of the features enabled.

This is an MPI application and should be run under `mpirun` or its equivalent.

The following example runs 16 processes and a specified hosts file using the default options (all-to-all connectivity, 64 to 4MB messages in powers of two, one iteration, no data integrity checking):

```
$ mpirun -np 16 -m hosts mpi_stress
```

There are a number of options for `mpi_stress`; this one may be particularly useful:

```
-P
```

This option *poisons* receive buffers at initialization and after each receive; pre-initialize with random data so that any parts that are not being correctly updated with received data can be observed later.

See the `mpi_stress(1)` man page for more information.

rpm

To check the contents of an installed RPM, use these commands:

```
$ rpm -qa infinipath\* mpi-\*
$ rpm -q --info infinipath # (etc)
```

The option `-q` queries. The option `--qa` queries all. To query a package that has not yet been installed, use the `-qpl` option.

strings

Use the `strings` command to determine the content of and extract text from a binary file.

The command `strings` can also be used. For example, the command:

```
$ strings -a /usr/lib/libinfinipath.so.4.0 | grep Date:
```

produces this output:

```
$Date: 2009-02-26 12:05 Release2.3 InfiniPath $
```

NOTE

The `strings` command is part of `binutils` (a development RPM), and may not be available on all machines.

Common Tasks and Commands

Table G-3 lists some common commands that help with administration and troubleshooting. Note that `mpirun` in `nonmpi` mode can perform a number of checks.

Table G-3. Common Tasks and Commands Summary

Function	Command

Check the system state	`ipath_checkout [options] hostsfile`
	`ipathbug-helper -m hostsfile \`
	`> ipath-info-allhosts`
	`mpirun -m hostsfile -ppn 1 \`
	`-np numhosts -nonmpi ipath_control -i`
	Also see the file:
	`/sys/class/infini-band/ipath*/device/status_str`
	where * is the unit number. This file provides information about the link state, possible cable/switch problems, and hardware errors.

Verify hosts via an Ethernet ping	`ipath_checkout --run=1 hostsfile`
Verify `ssh`	`ipath_checkout --run=2 hostsfile`
Show `uname -a` for all hosts	`mpirun -m hostsfile -ppn 1 \` `-np numhosts -nonmpi uname -a`

Reboot hosts	As a root user:
	`mpirun -m hostsfile -ppn 1 \`
	`-np numhosts -nonmpi reboot`

Run a command on all hosts	`mpirun -m hostsfile -ppn 1 \`
	`-np numhosts -nonmpi <command>`
	Examples:
	`mpirun -m hostsfile -ppn 1 \`
	`-np numhosts -nonmpi hostname`
	`mpirun -m hostsfile -ppn 1 \`
	`-np numhosts -nonmpi date`

Copy a file to all hosts	Using bash:
	`$ for i in $( cat hostsfile )`
	`do`
	`scp <source> $i:<destination>`
	`done`

Summarize the fabric components	`ipathbug-helper -m hostsfile \` `> ipath-info-allhosts`

Table G-3. Common Tasks and Commands Summary (Continued)

Function	Command

Show the status of host IB ports	`ipathbug-helper -m hostsfile \`
	`> ipath-info-allhosts`
	`mpirun -m hostsfile -ppn 1 \`
	`-np numhosts -nonmpi ipath_control -i`

Verify that the hosts see each other	`ipath_checkout --run=5 hostsfile`
Check MPI performance	`ipath_checkout --run=7 hostsfile`
Generate all hosts problem report information	`ipathbug-helper -m hostsfile \` `> ipath-info-allhosts`

Table Notes

The " \ " indicates commands that are broken across multiple lines.

Summary and Descriptions of Useful Files

Useful files are summarized in Table G-4. Names in blue text are linked to a corresponding section that provides further details.

Table G-4. Useful Files

File Name	Function
`boardversion`	File that shows the version of the chip architecture.
`status_str`	File that verifies that the InfiniPath software is loaded and functioning
`/var/log/messages`	Logfile where various programs write messages. Tracks activity on your system
`version`	File that provides version information of installed software/drivers

boardversion

It is useful to keep track of the current version of the chip architecture. You can check the version by looking in this file:

```
/sys/class/infiniband/qib0/device/boardversion
```

Example contents are:

```
ChipABI 2.0,InfiniPath_QLE7280,InfiniPath1 5.2,PCI 2,SW
Compat 2
```

This information is useful for reporting problems to Technical Support.

NOTE

This file returns information of where the form factor adapter is installed. The PCIe half-height, short form factor is referred to as the QLE7140, QLE7240, QLE7280, QLE7340, or QLE7342.

status_str

Check the file `status_str` to verify that the InfiniPath software is loaded and functioning. The file is located here:

```
/sys/class/infiniband/qib/device/status_str
```

Table G-5 shows the possible contents of the file, with brief explanations of the entries.

Table G-5. `status_str` File Contents

File Contents	Description
`Initted`	The driver has loaded and successfully initialized the IBA6110 or IBA7220 ASIC.
`Present`	The IBA6110 or IBA7220 ASIC has been detected (but not initialized unless Initted is also present).
`IB_link_up`	The IB link has been configured and is in the active state; packets can be sent and received.
`IB_configured`	The IB link has been configured. It may or may not be up and usable.
`NOIBcable`	Unable to detect link present. This problem can be caused by one of the following problems with the QLE7140, QLE7240, or QLE7280 adapters: ■ No cable is plugged into the adapter. ■ The adapter is connected to something other than another IB device, or the connector is not fully seated. ■ The switch where the adapter is connected is down.
`Fatal_Hardware_Error`	Check the system log (default is `/var/log/messages`) for more information, then call Technical Support.

This same directory contains other files with information related to status. These files are summarized in Table G-6.

Table G-6. Status—Other Files

File Name	Contents
`lid`	IB LID. The address on the IB fabric, similar conceptually to an IP address for TCP/IP. *Local* refers to it being unique only within a single IB fabric.
`mlid`	The Multicast Local ID (MLID), for IB multicast. Used for InfiniPath ether broadcasts, since IB has no concept of broadcast.
`guid`	The GUID for the InfiniPath chip, it is equivalent to a MAC address.
`nguid`	The number of GUIDs that are used. If `nguids=2` and two chips are discovered, the first chip is assigned the requested GUID (from `eeprom`, or `ipath_sma`), and the second chip is assigned GUID+1.
`serial`	The serial number of the QLogic adapter.
`unit`	A unique number for each card or chip in a system.
`status`	The numeric version of the `status_str` file, described in Table G-5.

`version`

You can check the version of the installed InfiniPath software by looking in:

```
/sys/class/infiniband/qib0/device/driver/version
```

QLogic-built drivers have contents similar to:

```
$Id: QLogic OFED Release x.x.x$ $Date: Day mmm dd hh:mm:ss
timezone yyyy $
```

Non-QLogic-built drivers (in this case `kernel.org`) have contents similar to:

```
$Id: QLogic kernel.org driver $
```

Summary of Configuration Files

Table G-7 contains descriptions of the configuration and configuration template files used by the InfiniPath and OpenFabrics software.

Table G-7. Configuration Files

Configuration File Name	Description
/etc/modprobe.conf	Specifies options for modules when added or removed by the modprobe command. Also used for creating aliases. The PAT write-combing option is set here. For Red Hat 5.*X* systems.
/etc/modprobe.d/ib_qib.conf	Specifies options for modules when added or removed by the modprobe command. Also used for creating aliases. The PAT write-combing option is set here. For Red Hat 6.*X* systems.
/etc/modprobe.conf.local	Specifies options for modules when added or removed by the modprobe command. Also used for creating aliases. The PAT write-combing option is set here. For SLES systems.
/etc/infiniband/openib.conf	The primary configuration file for InfiniPath, OFED modules, and other modules and associated daemons. Automatically loads additional modules or changes IPoIB transport type.
/etc/sysconfig/infinipath	Contains settings, including the one that sets the ipath_mtrr script to run on reboot.
/etc/sysconfig/network/ifcfg-<*NAME*>	Network configuration file for network interfaces For SLES systems.
/etc/sysconfig/network-scripts/ifcfg-<*NAME*>	Network configuration file for network interfaces For Red Hat systems.
Sample and Template Files	**Description**
/usr/share/doc/initscripts-*/sysconfig.txt	File that explains many of the entries in the configuration files For Red Hat systems.

H Recommended Reading

Reference material for further reading is provided in this appendix.

References for MPI

The MPI Standard specification documents are located at:

http://www.mpi-forum.org/docs

The MPICH implementation of MPI and its documentation are located at:

http://www-unix.mcs.anl.gov/mpi/mpich/

The ROMIO distribution and its documentation are located at:

http://www.mcs.anl.gov/romio

Books for Learning MPI Programming

Gropp, William, Ewing Lusk, and Anthony Skjellum, *Using MPI*, Second Edition, 1999, MIT Press, ISBN 0-262-57134-X

Gropp, William, Ewing Lusk, and Anthony Skjellum, *Using MPI-2*, Second Edition, 1999, MIT Press, ISBN 0-262-57133-1

Pacheco, *Parallel Programming with MPI,* 1997, Morgan Kaufman Publishers, ISBN 1-55860

Reference and Source for SLURM

The open-source resource manager designed for Linux clusters is located at:

http://www.llnl.gov/linux/slurm/

InfiniBand®

The InfiniBand® specification can be found at the InfiniBand® Trade Association (IBTA) website:

http://www.infinibandta.org/

OpenFabrics

Information about the OpenFabrics Alliance (OFA) is located at:

http://www.openfabrics.org

Clusters

Gropp, William, Ewing Lusk, and Thomas Sterling, *Beowulf Cluster Computing with Linux*, Second Edition, 2003, MIT Press, ISBN 0-262-69292-9

Networking

The Internet Frequently Asked Questions (FAQ) archives contain an extensive Request for Command (RFC) section. Numerous documents on networking and configuration can be found at:

http://www.faqs.org/rfcs/index.html

Rocks

Extensive documentation on installing Rocks and custom Rolls can be found at:

http://www.stackiq.com

Other Software Packages

Environment Modules is a popular package to maintain multiple concurrent versions of software packages and is available from:

http://modules.sourceforge.net/

Corporate Headquarters QLogic Corporation 26650 Aliso Viejo Parkway Aliso Viejo, CA 92656 949.389.6000 **www.qlogic.com**

International Offices UK | Ireland | Germany | France | India | Japan | China | Hong Kong | Singapore | Taiwan

© 2012 QLogic Corporation. Specifications are subject to change without notice. All rights reserved worldwide. QLogic, the QLogic logo, and the Powered by QLogic logo are registered trademarks of QLogic Corporation. InfiniBand is a registered trademark of the InfiniBand Trade Association. All other brand and product names are trademarks or registered trademarks of their respective owners. Information supplied by QLogic Corporation is believed to be accurate and reliable. QLogic Corporation assumes no responsibility for any errors in this brochure. QLogic Corporation reserves the right, without notice, to make changes in product design or specifications.